

ZOOMING IN ON THE MULTI-PHASE STRUCTURE OF MAGNETICALLY-DOMINATED QUASAR DISKS: RADIATION FROM TORUS TO ISCO ACROSS ACCRETION RATES

PHILIP F. HOPKINS^{1,*}, KUNG-YI SU², NORMAN MURRAY³, ULRICH P. STEINWANDEL⁴, NICHOLAS KAAZ⁵, SAM B. PONNADA¹, JAEDEN BARDATI¹, JOANNA M. PIOTROWSKA¹, HAI-YANG WANG¹, YANLONG SHI³, DANIEL ANGLÉS-ALCÁZAR⁶, ELIAS R. MOST¹, KYLE KREMER¹, CLAUDE-ANDRÉ FAUCHER-GIGUÈRE⁶, AND SARAH WELLONS⁷

¹ TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

² Black Hole Initiative, Harvard University, 20 Garden St, Cambridge, MA 02138, USA

³ Canadian Institute for Theoretical Astrophysics, University of Toronto, Toronto, ON M5S 3H8, Canada

⁴ Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave., New York, NY 10010 USA

⁵ CIERA and Department of Physics and Astronomy, Northwestern University, 1800 Sherman Ave, Evanston, IL 60201, USA

⁶ Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269, USA and

⁷ Department of Astronomy, Van Vleck Observatory, Wesleyan University, 96 Foss Hill Drive, Middletown, CT 06459, USA

Version April 29, 2025

ABSTRACT

Recent radiation-thermochemical-magnetohydrodynamic simulations resolved formation of quasar accretion disks from cosmological scales down to ~ 300 gravitational radii R_g , arguing they were “hyper-magnetized” ($\text{plasma } \beta \ll 1$ supported by toroidal magnetic fields) and distinct from traditional α -disks. We extend these, refining to $\approx 3 R_g$ around a $10^7 M_\odot$ BH with multi-channel radiation and thermochemistry, and exploring a factor of 1000 range of accretion rates ($\dot{m} \sim 0.01 - 20$). At smaller scales (depending on \dot{m}), we see the disks maintain steady accretion, thermalize and self-ionize, and radiation pressure grows in importance, but large deviations from local thermodynamic equilibrium and single-phase equations of state are always present. Trans-Alfvénic and highly-supersonic turbulence persists in all cases, and leads to efficient vertical mixing, so radiation pressure saturates at levels comparable to fluctuating magnetic and turbulent pressures even for $\dot{m} \gg 1$. The disks also become radiatively inefficient in the inner regions at high \dot{m} . The midplane magnetic field remains primarily toroidal at large radii, but at super-Eddington \dot{m} we see occasional transitions to a poloidal-field dominated state associated with outflows and flares. Large-scale magnetocentrifugal and continuum radiation-pressure-driven outflows are weak at $\dot{m} < 1$, but can be strong at $\dot{m} \gtrsim 1$. In all cases there is a scattering “photosphere” above the disk extending to $\gtrsim 1000 R_g$ at large \dot{m} , and the disk is thick and flared owing to magnetic support (with H/R nearly independent of \dot{m}), so the outer disk is strongly illuminated by the inner disk and most of the inner disk continuum scatters or is reprocessed at larger scales, giving apparent emission region sizes as large as $\gtrsim 10^{16} \text{ cm}$.

Subject headings: quasars: general — accretion, accretion disks — quasars: supermassive black holes — galaxies: active — galaxies: evolution — galaxies: formation

1. INTRODUCTION

Accretion disks around supermassive black holes (SMBHs) are the engines that power active galactic nuclei (AGN) and quasars, the most luminous sustained sources known (Schmidt 1963; Salpeter 1964). Mass inflow rates exceeding $\gtrsim 1 M_\odot \text{ yr}^{-1}$ are required to explain typical quasars at redshifts $z \gtrsim 2$ (Shen et al. 2020), and to provide most of the SMBH mass today (Soltan 1982). The effects of “AGN feedback” in the form of radiation, outflows, and jets launched from such disks near the BH (Laor et al. 1997; Crenshaw et al. 2000; Dunn et al. 2010; Sturm et al. 2011; Faucher-Giguère & Quataert 2012; Faucher-Giguère et al. 2012; Zakamska et al. 2016; Williams et al. 2017) can also have dramatic effects (Silk & Rees 1998; King 2003; Di Matteo et al. 2005; Murray et al. 2005; Hopkins et al. 2005a,b; Torrey et al. 2020) potentially explaining BH-host galaxy correlations and co-evolution (Magorrian et al. 1998; Ferrarese & Merritt 2000; Gebhardt et al. 2000; Hopkins et al. 2007; Aller & Richstone 2007; Kormendy et al. 2011) and shaping galaxy “quenching” and masses (Croton et al. 2006; Hopkins et al. 2006a, 2008; Wellons et al. 2023; Mercedes-Feliz et al. 2023; Cochrane et al. 2023).

The theoretical study and modeling of accretion disks is

a rich field, and it has been recognized for decades that the structure and basic qualitative nature of the accretion disk, let alone structures surrounding or part of the outer disk such as the broad line region and dusty “torus,” are sensitive to the outer boundary conditions “supplying” the disk (Pringle 1981; Frank et al. 2002; Abramowicz & Fragile 2013; Jafari 2019). These flows come from interstellar medium (ISM) scales $\gg \text{pc}$ where star formation, non-equilibrium atomic+molecular radiation-thermochemistry, optically-thin cooling, and stellar feedback processes such as supernovae and stellar radiation must be modeled. In Hopkins et al. (2024b) (henceforth Paper I) we presented the first radiation-magnetohydrodynamic simulations to follow all of these physical effects in a single simulation beginning from a cosmological simulation with box size $\sim 100 \text{ Mpc}$ and refining down to scales of $\sim 80 \text{ au}$ or $\sim 300 R_g$,¹ and showed that indeed a clearly-defined, steady-state accretion disk forms interior to the BH radius of influence ($R_{\text{BHROI}} \sim 5 \text{ pc}$, interior to which the BH dominates the gravitational potential). As shown in Fig. 1, this particular event occurs in a high-redshift ($z \sim 4.5$) massive starburst (star formation rate $\gtrsim 200 M_\odot \text{ yr}^{-1}$) galaxy which is clumpy and merging, when a massive ($\gg 10^8 M_\odot$) star-forming cloud complex passes by and is tidally disrupted

¹ Here we define the gravitational radius equal to the Schwarzschild radius $R_g \equiv 2 G M_{\text{BH}}/c^2$, or $\approx 3.9 \times 10^{12} \text{ cm}$ here.

*E-mail: phopkins@caltech.edu

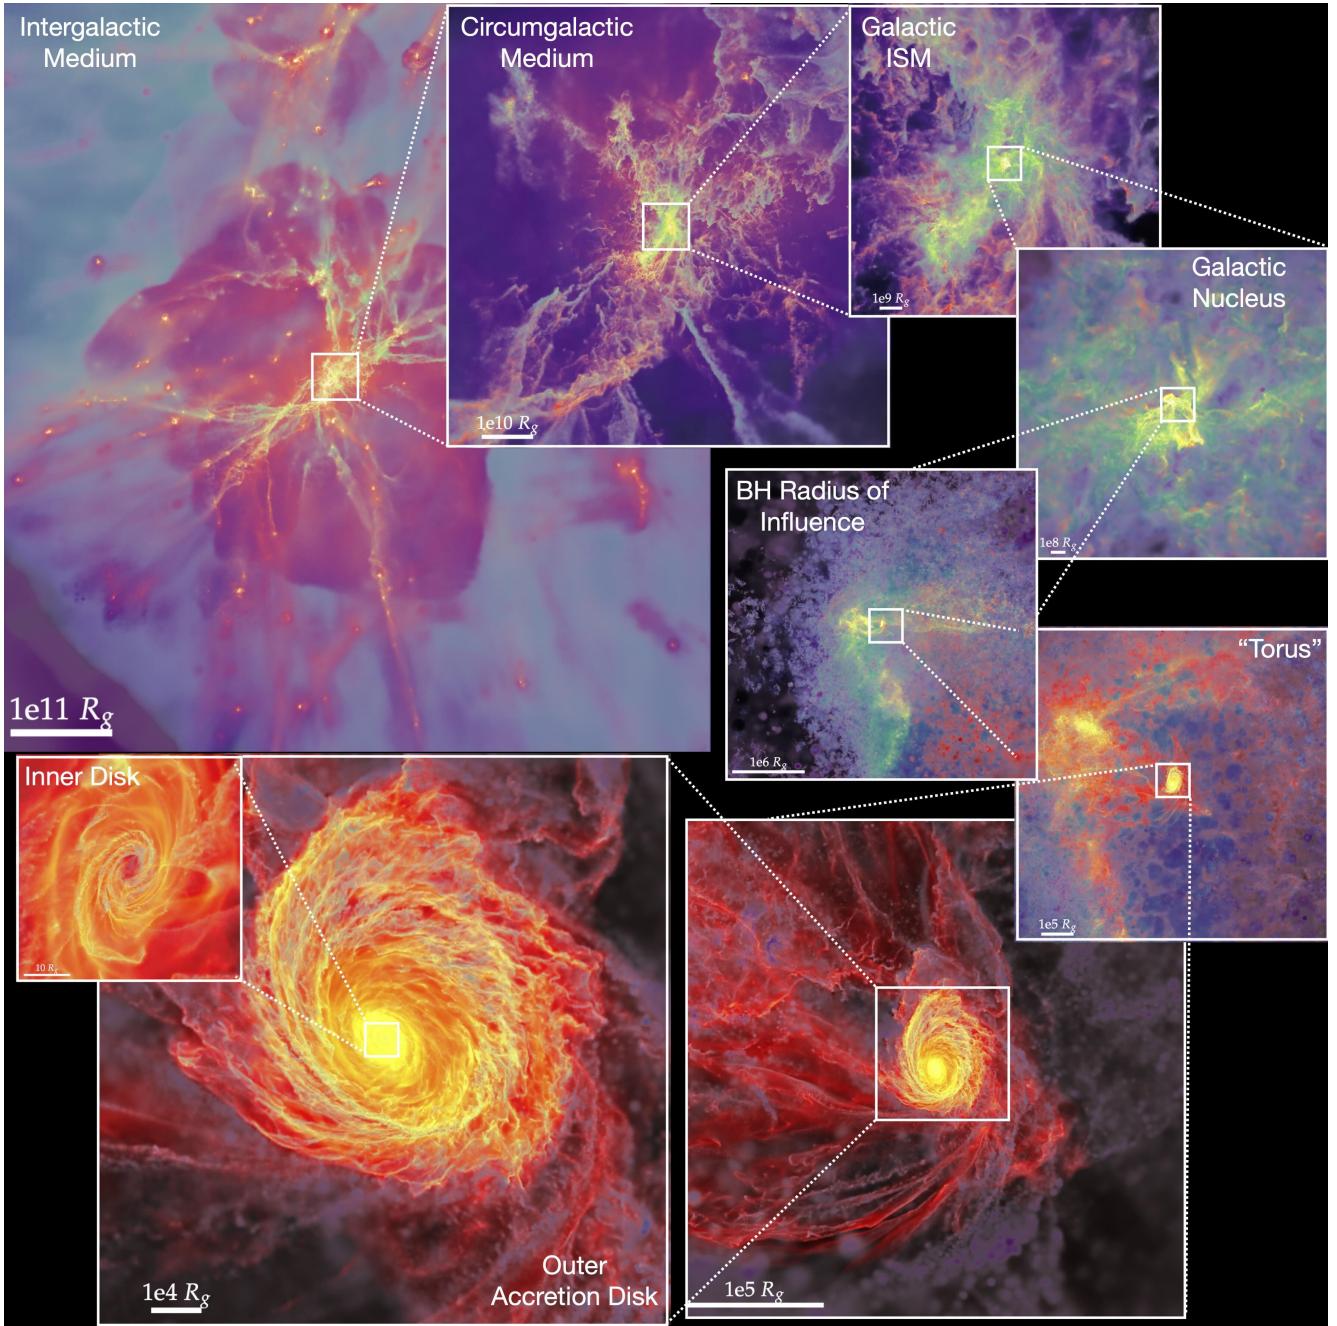


FIG. 1.— Images of the simulation on various scales. Bar shows the scale in units of $R_g \equiv 2GM_{\text{BH}}/c^2$ for the $1.3 \times 10^7 M_{\odot}$ BH at the galaxy center. Intensity scales with log of projected surface density (factor of $\sim 10^4$ range in each) and colormaps reflect the dominant pressure source (thermal plus radiation in purple-magenta, turbulent/kinetic in blue-green, magnetic in red-yellow). The simulation from Paper I begins from a fully-cosmological box and trace $\sim 10^{-15}$ G magnetic fields at redshift $z \approx 150$, zooming in to the formation of a massive starburst galaxy, which undergoes multiple mergers and gas-rich accretion at $z \sim 4.5$ (galaxy-integrated star formation rate $\sim 200 M_{\odot} \text{ yr}^{-1}$), wherein a massive ($\sim 10^8 M_{\odot}$) molecular cloud-complex is disrupted around the BH radius of influence ($\sim 5 \text{ pc}$ here) and forms the accretion disk. Here we continue to zoom in to the inner disk, seeing the accretion disk persist down to $\sim 3 R_g$.

by the pre-existing SMBH (mass $\sim 1.3 \times 10^7 M_{\odot}$), fueling up to $\sim 20 - 30 M_{\odot} \text{ yr}^{-1}$ steady-state inflows which fall into the BHROI and circularize forming the accretion disk at $\sim 0.1 - 1 \text{ pc}$.

In Hopkins et al. (2024c) (Paper II), we studied the detailed properties of these disks, showing that they were qualitatively unlike classic thermal-pressure-dominated α -disk models like that of Shakura & Sunyaev (1973) (hereafter SS73), nor did they resemble magnetically “arrested” or “elevated”

disks (Bisnovatyi-Kogan & Ruzmaikin 1976; Narayan et al. 2003). Instead, these disks are supported by strong magnetic fields with plasma $\beta \equiv \beta_{\text{thermal}} \equiv c_s^2/v_A^2 \sim P_{\text{gas, thermal}}/P_B \sim 10^{-6} - 10^{-2} \ll 1$ (see e.g. Gaburov et al. 2012),² where the field is dominated by the mean toroidal component and

² We adopt $c_s \equiv \sqrt{k_B T_{\text{gas}}/\mu m_p}$ throughout, i.e. the sound speed refers specifically the isothermal gas sound speed (as opposed to some “effective” wavespeed), for the reasons shown explicitly in § 4.9 where we discuss the degree to which strict tight-coupling of radiation is inapplicable here.

secondary radial and turbulent/fluctuating components (with weak, turbulently-sourced mean poloidal field), amplified primarily via flux-freezing from modest/typical ISM fields. They are also trans-Alfvénically turbulent, which necessarily means the turbulence is highly supersonic and that the medium has cooling times t_{cool} much smaller than dynamical times $t_{\text{dyn}} \sim \Omega^{-1} \equiv (G M_{\text{BH}}/R^3)^{-1/2}$ (with cooling balancing heating). In Hopkins et al. (2024a), we studied how star formation is completely suppressed in such disks, noting that they resolve the decades-old problem that thermal-pressure dominated disks should be gravitationally unstable (at high accretion rates) outside a few dozen gravitational radii. And in Hopkins et al. (2024d) (Paper III) we presented some simple analytic similarity models which describe the simulation outer disks and attempted to extrapolate them to smaller radii. Since then, similar disks have been seen under very different conditions, using different codes and numerical methods to simulate accretion from large-to-small scales in massive ellipticals (at much lower accretion rates; Guo et al. 2024) and from star clusters onto intermediate-mass BHs (Shi et al. 2024a). This suggests they may be ubiquitous in AGN.

However, these studies leave a number of behaviors at smaller radii uncertain. Most notably, the particular accretion episode studied in Paper I-Paper II features a supercritical Eddington-scaled accretion rate $\dot{m} \equiv \dot{M}/(M_{\text{BH}}/t_S) = 0.1 \dot{M} c^2/L_{\text{Edd}} \sim 100$ (with $t_S \approx 5 \times 10^7$ yr). But with the simulation only reaching down to $\sim 300 R_g$, the actual accretion luminosity in the *resolved* portion of the disk is everywhere sub-Eddington. Based on the analytic extrapolations in Paper III, much of the interesting behavior would be expected to appear at the next order of magnitude down in scale, $\sim 30 - 300 R_g$. At these radii, for a simple radiatively-efficient model akin to SS73, we expect an order-unity fraction of the total luminosity to be produced. And somewhere in this range, if $\dot{m} \gg 1$ were maintained, models predict that the radiation pressure becomes comparable to (or larger than) the magnetic pressure. Moreover, if the accretion disk is effectively optically thick and radiation transport were strictly diffusive, so the midplane temperature $T_{\text{mid}} \sim T_{\text{eff}} \tau^{1/4}$, then simple analytic arguments suggest any disk would have radiation pressure internal to the disk sufficient to inflate the disk to $H \gtrsim R$ and transition to either (a) radiatively inefficient (slim disk-like) accretion or (b) very strong outflows (greatly decreasing \dot{m}), at radii $x_g \equiv R/R_g \gtrsim 3 \dot{m}$. But this is just interior to the boundary in Paper II. And as noted in Paper III, the strong turbulence in these disks could *qualitatively* change this conclusion, with major implications for both the structure of magnetically-dominated disks (at high- \dot{m}) and the possibility of super-critical accretion.

A number of other important physical transitions are predicted to occur, based on the extrapolations in Paper III, at scales $\lesssim 1000 R_g$. Most notably, within this range the radiation should become hard enough that the disk fully self-ionizes (it is primarily atomic at $\gtrsim 1000 R_g$), and the scattering+absorption optical depth should become large enough that the “effective” optical depth $\sim \sqrt{\kappa_a(\kappa_a + \kappa_s)} \Sigma_{\text{gas}}$ exceeds unity and the disk could thermalize with optically-thick quasi-blackbody cooling. So in this range it may still be important to include physics from Paper I-Paper II usually neglected in accretion disk simulations such as detailed non-equilibrium atomic/molecular chemistry, optically-thin plus thick multi-group radiation and coupled heating/cooling (including recombination, H^- , H , He and metal line, as well as continuum

processes), and perhaps even self-gravity. And the study in Paper I-Paper II only considered one accretion rate, so it is challenging to make extrapolations or test different analytic models for the disk structure.

In this paper, we therefore extend the simulations presented in Paper I-Paper II by further refining down to $\sim 3 R_g$ around the SMBH, while also modifying the outer boundaries to produce a range of accretion rates. This allows us to capture the key transitions in the disk described above as a function of both radius and accretion rate, before reaching radii where the physics limitations of the simulations – we do not include general relativity [GR] nor explicit treatment of X-ray radiation/ionization nor weakly-collisional two-temperature plasmas – become problematic. In § 2 we summarize our numerical methods and different tests performed (reviewing the basic physics § 2.1, initial conditions § 2.2, and accretion rate survey § 2.3). In § 3 we present and discuss the results, including convergence and steady-state behaviors (§ 3.1), results similar to those at larger radii (§ 3.2) and unique to smaller radii (§ 3.3), luminosities and radiative efficiencies (§ 3.4), and variability/bursts/state changes (§ 3.5). In § 4 we discuss the physics of these results, attempting to understand the radiation pressure and its saturation (§ 4.1), radiative efficiencies (§ 4.2), scaling with accretion rate \dot{m} (§ 4.3), comparison to analytic similarity models (§ 4.4), physics of extended scattering/reprocessing regions (§ 4.5) and implications for the emergent spectrum (§ 4.6), properties of outflows (§ 4.7) and of the broad-line region (§ 4.8), deviations from local thermodynamic equilibrium and simple equation-of-state models (§ 4.9), and sensitivity to different physics (§ 4.10). We summarize and conclude in § 5. In the Appendices we include more details of the cooling physics and opacities (§ A) and the full set of physical and numerical variations considered (§ B), discuss the effective spatial resolution of the simulations (§ C) and show tests where the thermal scale-length is always forced to be well-resolved (§ D).

2. SIMULATIONS & METHODS

2.1. Overview & Basic Physics

The simulations here are direct continuations/extensions of those presented in Paper I, to which we refer for details. In brief, we initialized a $\sim (100 \text{ cMpc})^3$ cosmological periodic box at redshift $z \sim 100$, with a trace primordial magnetic field ($\sim 10^{-15} \text{ G}$), and follow it as a cosmological galaxy formation simulation using the full combined Feedback In Realistic Environments (FIRE) and STARFORGE physics (Hopkins et al. 2018, 2023a; Grudić et al. 2021; Guszejnov et al. 2021), with initial “coarse” adaptive resolution reaching $\sim 10 \text{ pc}$ in the galactic nucleus. At redshift $z \sim 4.5$ the galaxy is a massive, high-redshift starburst (star formation rate $\gtrsim 200 \text{ M}_\odot \text{ yr}^{-1}$) and a massive molecular cloud complex is perturbed to encounter the supermassive black hole of mass $\sim 1.3 \times 10^7 \text{ M}_\odot$. We therefore hyper-refine at this time to follow the evolution in the central $\lesssim 10 - 100 \text{ pc}$ of the galaxy. In Paper I we refine to a uniform target resolution of $\Delta m \sim 0.005 \text{ M}_\odot$ in the central $\lesssim 5 \text{ pc}$ (the BH radius of influence, BHROI, where it begins to dominate the total enclosed mass and gravitational potential), following the gas inflow from the tidally disrupted GMC complex. This material circularizes at $\sim 0.1 - 1 \text{ pc}$, forming a stable, non-star-forming, magnetically-dominated (Paper II) accretion disk around the BH, which we evolved in Paper I-Paper III down to an inner boundary at 80 au $\approx 300 R_g$ from the BH (where $R_g \equiv 2 G M_{\text{BH}}/c^2$ is the Schwarzschild

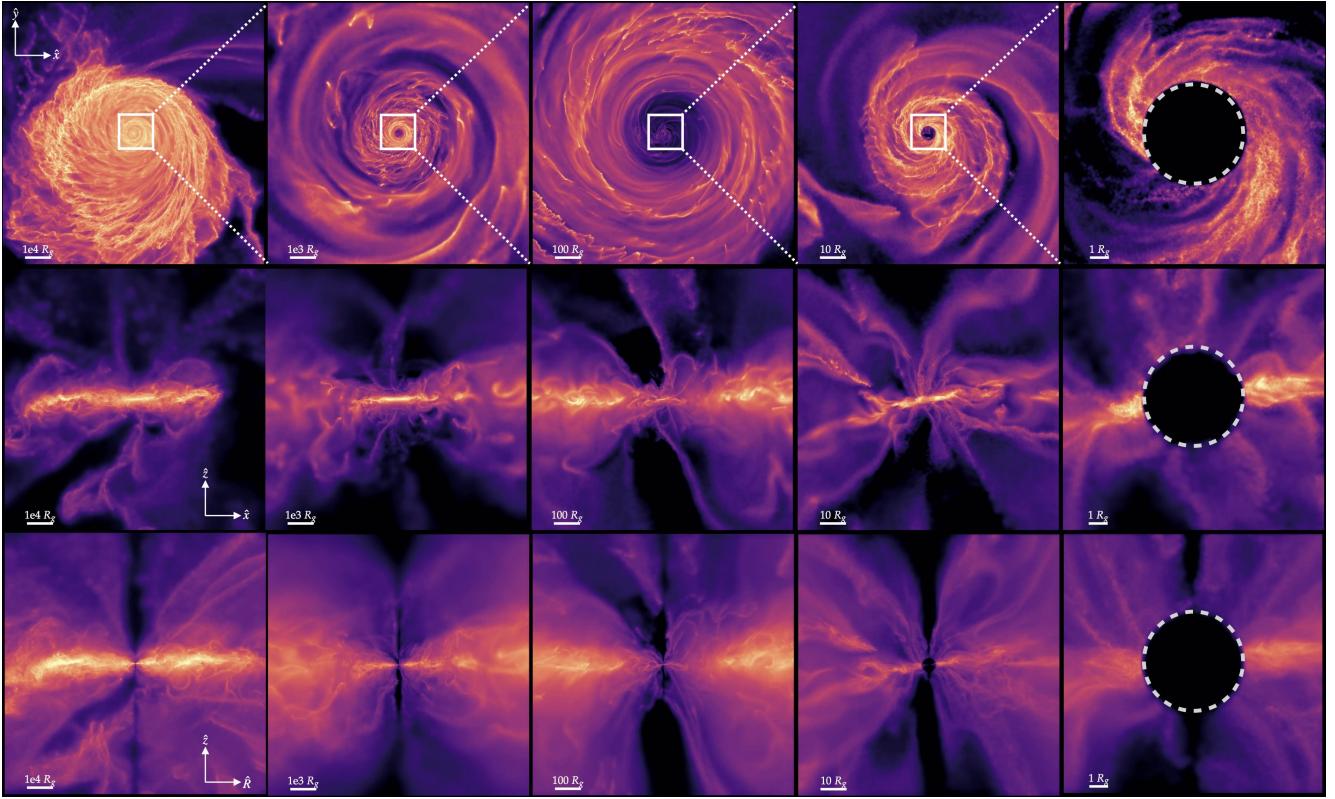


FIG. 2.— Surface density projection (increasing black-to-yellow, scaled independently in each panel to show the 1 – 99% range in each image) of the accretion disk at various scales from $\sim 10^5 R_g \sim 0.1$ pc (the outer circularization radius) to $\sim 3 R_g$ (the ISCO and inner accretion boundary of the simulation; *dashed circle*), with scales labeled (left-to-right). Rows show different projections: face-on slice through the x - y plane (*top*), edge-on slice through the x - z plane (*middle*), and edge-on cylindrical R - z projection (in a wedge of azimuthal opening angle $\sim 30^\circ$). The accretion flow is clearly disk-like, but with a thick, turbulent structure with material at large H/R at all radii here.

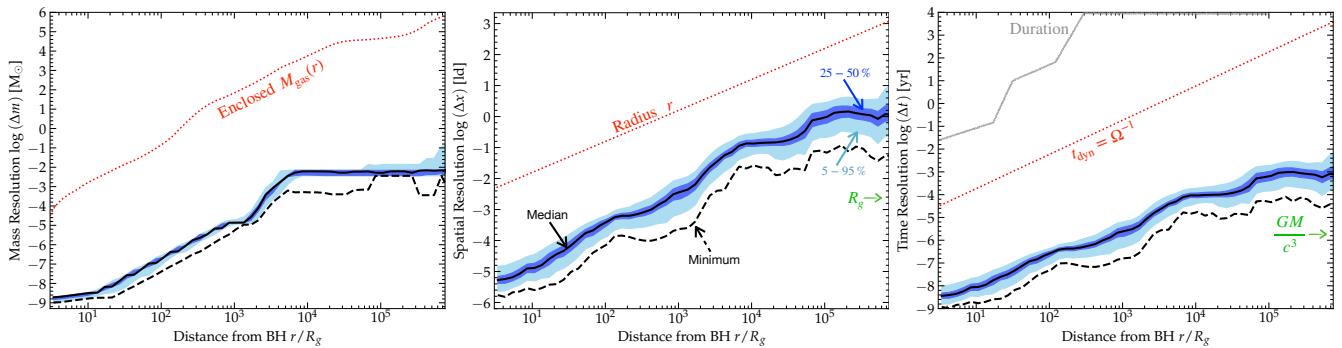


FIG. 3.— Resolution of our simulations as a function of distance from the BH r (in units of R_g), in the accretion disk from the ISCO at $\sim 3 R_g$ to beyond the circularization radius at ~ 1 pc ($\sim 10^6 R_g$). In spherical shells, we plot the gas-mass-weighted median (*solid*) and minimum (*dashed*) cell mass Δm (*left*), linear size $\Delta x \equiv (\Delta m/\rho)^{1/3}$ (*center*), and numerical timestep Δt (*right*), with their interquartile and 90% range. We compare the total enclosed disk mass $M_{\text{gas}}(< r)$, size scale r , and dynamical time $t_{\text{dyn}} \equiv 1/\Omega \approx \sqrt{r^3/GM_{\text{BH}}}$, as well as the total duration/run-time of the simulations after reaching the maximum refinement level at each radius. We also note the value of $R_g \equiv 2GM_{\text{BH}}/c^2$ and GM_{BH}/c^3 . In the midplane of the disk region we study here, $\Delta x \sim 0.0001 - 0.001 R$ (akin to a $1000^3 - 10,000^3$ simulation at each log radius), with the smaller value representative of the resolution in dense substructures and the larger representative of the lower-density material in the disk.

radius). Here we extend this refinement; the range of scales is illustrated in Fig. 1.

The simulations include magnetic fields (using the hybrid Powell et al. 1999; Dedner et al. 2002 plus constrained-gradient method from Hopkins & Raives 2016; Hopkins 2016), with anisotropic Braginskii viscosity and conduction (Su et al. 2017; Hopkins 2017); a variable local cosmic-ray background (Hopkins et al. 2022a; Hopkins 2023; Hopkins et al. 2022b,c) here modeled using the simple sub-grid

method from Hopkins et al. (2023b); self-gravity with adaptive self-consistent softenings scaling with the resolution and high-order Hermite integrators capable of accurately integrating $\gtrsim 10^5$ orbits in hard binaries (Grudić & Hopkins 2020; Grudić et al. 2021; Grudić 2021; Hopkins et al. 2023c); metal enrichment and dust destruction/sublimation (Ma et al. 2017; Gandhi et al. 2022; Choban et al. 2022); super-massive black hole seed formation and growth via gravitational capture of gas (Hopkins et al. 2016; Shi et al. 2023; Wellons et al. 2023);

(proto)star formation and accretion and explicit feedback from stars in the form of protostellar jets, main-sequence stellar mass-loss, radiation, and supernovae (Grudić et al. 2022; Guszejnov et al. 2022a,b, 2023). The inner boundary is an inflow/sink boundary around the BH. The simulations evolve explicit multi-group radiation-hydrodynamics using the M1 approximation (Levermore 1984) with adaptive-wavelength bands (Hopkins et al. 2020; Hopkins & Grudić 2019; Grudić et al. 2021) coupled explicitly to all the thermo-chemical processes, together with radiative cooling and thermo-chemistry incorporating cosmic backgrounds, radiation from local stars, re-radiated cooling radiation, dust, molecular, atomic, metal-line, and ionized species opacities, cosmic rays, and other processes. This allows us to self-consistently model the thermochemistry and opacities in gas with densities from densities $n \sim 10^{-8} - 10^{16} \text{ cm}^{-3}$ and temperatures $\sim 1 - 10^{10} \text{ K}$ in a range of radiation and cosmic ray environments (see Paper I).

Our default RHD treatment includes several different evolved radiation groups: OIR, NUV, EUV, and ionizing (H, He) bands with fixed photon frequency boundaries at $\sim (0.41, 3.4, 8, 13.6, 24.6, 100) \text{ eV}$, plus an adaptive self-radiation band which separately evolves the re-emitted continuum radiation from gas and/or dust as a greybody-type spectral shape but with explicitly-evolved energy and photon number density, e.g. explicitly evolving its radiation temperature $T_{\text{rad}}(\mathbf{x}, t)$ or effective frequency $\langle v_{\text{rad}} \rangle(\mathbf{x}, t)$ independently in each cell, rather than assuming local thermodynamic equilibrium (LTE) or strictly blackbody/thermal radiation. The latter can contribute radiation to all of the fixed-wavelength narrow bands, if T_{rad} is hot enough (as it is here), and all the relevant opacities and thermochemical couplings are appropriately Rosseland or Planck integrated for the evolved $T_{\text{rad}}(\mathbf{x}, t)$ (shape/frequency of the band, as well as gas properties like T_{gas} , density, metallicity, ionization state, etc.; see Appendix A). We also specifically evolve the coupled anisotropic photon-gas momentum equations with terms up to and including $\mathcal{O}(v_{\text{gas}}^2/c^2)$ (Mihalas & Mihalas 1984), rather than assuming the common “trapped photon” limit (which treats the photons as an isotropic “pressure” in the Riemann problem). Details of the opacities and radiation-thermochemical coupling relevant on these scales are given in Appendix A. The narrow fixed bands are designed primarily for situations like ISM chemistry where photo-electric heating, Lyman-Werner radiation, H ionization, and other processes are primarily sensitive to photons in a narrow frequency range and dominated by light emitted by stars, not the cooling radiation of gas. So the latter (adaptive) band contains most of the radiation energy density at the radii of interest here, with $T_{\text{rad}}(\mathbf{x}, t)$ as we show ranging from $\lesssim 30 \text{ K}$ to $\gtrsim 10^8 \text{ K}$. This allows us to flexibly handle continuum opacities over an extremely broad range of conditions without making assumptions about a single grey opacity or even source of opacities, whether the system is in LTE or non-LTE, optically thick or thin, etc. However, we caution that our treatment does not include resonant line scattering, so we cannot self-consistently capture line-driven winds.

In this paper we extend these simulations by further refining in the inner regions, shrinking our inner sink boundary in a couple of stages, first from $\sim 300 \rightarrow 30 R_g$ (with much higher gas resolution in the inner regions), then further to $\rightarrow 3 R_g$ (with a Paczyński-Wiita [PW] BH potential). Because we are simulating even smaller radii, we also remove the “reduced speed of light” approximation used in Paper I for the radiation-

hydrodynamics. There, a reduced $\tilde{c} \approx 0.1 c$ was used; here given the large speeds reached in the center we simply adopt $\tilde{c} = c$, i.e. no reduction.³ We retain all of the other physics above. Fig. 2 illustrates this higher-resolution region.

Appendix B describes a number of experiments with numerical and physical parameters which we ran for much shorter times, including different refinement schemes, resolution variations, varying the inner/outer boundary conditions, re-initializing some quantities in the initial conditions, removing star formation and stellar evolution physics, switching between PW and Keplerian potentials, and experimenting with variations to the detailed cooling-radiation-thermochemical coupling and approximations. These generally have no impact on our qualitative conclusions, except where we otherwise specify.

2.2. Initial Conditions and Resolution

Specifically, for our fiducial simulation here, we restart from the final snapshot of the Paper I-Paper II simulations described above (which had been evolved for $\sim 10^5$ dynamical times at the inner radius of $\sim 300 R_g$), and refine further. We initially reduce the size of the inner accretion boundary to $30 R_g$, and we turn off the speed of light reduction throughout. We apply an additional refinement layer, enforcing a target mass resolution of $\Delta m \propto (r/0.01 \text{ pc})^4$ in a boundary layer between $\sim 0.001 - 0.02 \text{ pc}$ (exterior to which the rest of the accretion disk has target resolution $\Delta m \approx 0.005 M_\odot$) to a minimum and approximately uniform mass resolution of $\Delta m \approx 3 \times 10^{-7} M_\odot$ interior to $< 0.001 \text{ pc}$ ($\sim 200 \text{ au}$ or $\sim 800 R_g$). This initial uniform-resolution high-resolution region contains approximately $\approx 10^7$ roughly equal-mass cells, which then increases to our target resolution after the simulation turns on (see Fig. 3). The “outer boundary” of our highest-resolution region of this simulation is therefore at $\sim 1000 R_g$, and the outer boundary of the intermediate resolution region (the high-resolution region from Paper I) is at $\sim 4 \times 10^6 R_g$ ($\sim 5 \text{ pc}$), but our box extends to $\sim 100 \text{ Mpc}$ ($10^{14} R_g$). After $\sim 10 \text{ yr}$ physical time of evolution ($\sim 10^4 \Omega^{-1}$ at this inner boundary of $\sim 30 R_g$) we iterate, moving the inner boundary to $\sim 3 R_g$ and applying one further refinement layer with the target resolution decreasing as $\propto r^2$ interior to $< 0.001 \text{ pc}$ to a uniform $\Delta m \sim 10^{-9} M_\odot$ in the inner disk.⁴ We evolve this highest-resolution setup for a time $\sim 10^5 G M_{\text{BH}}/c^3$, i.e. $\sim 0.2 \text{ yr}$ (1800 hours) physical time or $\sim 7000 \Omega^{-1}$ at the inner boundary ($3 R_g$) and $\sim 200 \Omega^{-1}$ at the previous inner boundary ($\sim 30 R_g$). For the BH specifically, we replace its Keplerian acceleration in the gravity tree with that of a Paczyński-Wiita potential ($G M_{\text{BH}}/(r - R_g)$), though for the run with innermost radius of $\sim 30 R_g$, this is never more than a few percent-level correction and our tests below suggest it is largely negligible, but upon going to $\sim 3 R_g$ it will be more important. For the self-gravity of the gas (which must still obey the Poisson equation) we retain Newtonian gravity.

The vast majority of the material at the large radii in the simulations will do nothing (at resolved scales) over the duration of our small-scale simulation here. So in this simulation, for numerical convenience, we effectively “freeze” all the matter

³ Note this has no immediate effect on evolved conserved quantities, such as radiation energy densities, only on the dynamical equations used to evolve the radiation.

⁴ We also enforce a secondary, but less important in practice, spatial refinement criterion requiring the cell size be $< 0.03 r$ at all r in the high-resolution region.

at radii $r > 1$ pc ($\sim 10^6 R_g$), for memory/storage purposes saving it in a separate snapshot and simply adding the potential from this material analytically to the “live” material at small radii. The dynamical time $t_{\text{dyn}} = \Omega^{-1}$ (with $\Omega \approx \sqrt{G M_{\text{BH}}/r^3}$) at our innermost boundary is $\approx 3 \times 10^4$ s ~ 8 hours at $30 R_g$ and ≈ 1000 s at $3 R_g$, but at this “frozen” boundary it is $\approx 1.3 \times 10^{11}$ s (~ 4000 yr), so this is a reasonable approximation. It is more useful to think of the material at $\sim 0.01 - 1$ pc in our simulation (where the total time we evolve is $\lesssim 1$ orbital time) as forming our outer boundary condition.

Note that for a standard α -disk (SS73) with the dissipated energy scaling as $\propto \dot{M} \Omega^2 (1 - \sqrt{3} R_g/R)$ around a non-rotating BH, we would expect $\sim 25\%$ ($\sim 100\%$) of the total integrated bolometric luminosity to be released at $\gtrsim 30 R_g$ ($\gtrsim 3 R_g$). Appendix B describes a variety of experiments detailed below, where we modify the refinement scheme or physics, and run for much shorter times, and tests where we re-run from a different “initial” snapshots (to see if a different inner disk begins to form). These do not change our conclusions except where we explicitly note below.

2.3. Different Accretion Rates

Note that for this BH mass, we define the critical/Eddington accretion rate for a reference radiative efficiency = 0.1, $\dot{M}_{\text{crit}} \equiv L_{\text{Edd}}/0.1 c^2 \sim 0.3 M_{\odot} \text{ yr}^{-1}$, so $\dot{m} \equiv \dot{M}/\dot{M}_{\text{crit}}$.

The “default” initial condition in § 2.2 begins from the supercritical accretion rate at the outer boundaries predicted from our original galaxy-scale simulation. Here we wish to explore the effects of different accretion rates, in a controlled manner. One possibility is to repeat the zoom-in exercise from Paper I at different cosmic times or in different galaxies, beginning again from large-volume cosmological simulations (see [Anglés-Alcázar et al. 2021](#)). While such experiments are valuable and will be the subject of future work, they are extremely computationally expensive, and do not provide a controlled comparison, since everything else about the system varies as well. To instead enable a more straightforward comparison here, we adopt a simpler approach of rescaling the initial conditions in § 2.3. Specifically, since we are focused only on radii well interior to the BHROI, we can rescale the expected accretion rate by a factor $f_{\dot{M}}$, taking our initial condition (fiducial snapshot from Paper II) and multiplying all gas cell masses by $f_{\dot{M}}$ (keeping BH and other masses fixed), $\Delta m \rightarrow f \Delta m$. We rescale our “target” mass resolution in § 2.2 by the same factor. Following the analytic similarity models in Paper III, we rescale the initial magnetic fields $\mathbf{B} \rightarrow f^{1/2} \mathbf{B}$ and temperatures $T \rightarrow f T$ such that the Alfvén and sonic speeds v_A/v_K , c_s/v_K remain constant. In practice we find the temperature rescaling is unnecessary, since the disks have cooling times much shorter than dynamical times they almost immediately come into their new temperature equilibria. However, some \mathbf{B} rescaling is necessary, since the default initial conditions already feature large $v_A/v_K \sim H/R \sim 0.1 - 1$, so rescaling the gas mass by $f_{\dot{M}} \ll 1$ without rescaling \mathbf{B} would produce initial conditions with the gas everywhere unbound in the outer disk.

We emphasize that this rescaling is applied only once to initial conditions (snapshot of the outer disk with inner truncation radius $\sim 300 R_g$ from Paper II). We then repeat our refinement and simulation identically as in § 2.2, just with the target mass resolution shifted by the same factor so the spatial resolution, or number of cells in the disk, is roughly equal across all simulations. So this is akin to standard accretion disk studies using different initial conditions to represent different \dot{m} . Below and

in Appendix B we describe a number of detailed variants to how we initialize these disks in detail, which do not change our results.

In addition to our default initial condition ($f_{\dot{M}} = 1$), we simulate two values of $f_{\dot{M}}$: $f_{\dot{M}} = 0.03$ and $= 0.001$, which non-linearly give rise to accretion rates $\dot{m} \sim 0.3 - 0.4$ and $\sim 0.02 - 0.04$ ($\dot{M} \sim 0.1$ and $\sim 0.007 M_{\odot} \text{ yr}^{-1}$), respectively.

3. RESULTS

Fig. 1-2 illustrate the morphology and gas structure on a wide range of scales, though we will focus primarily on the smaller scales illustrated in Fig. 2 which are newly-resolved here. We first focus on a case study of our default ($f_{\dot{M}} = 1$) super-critical simulation, before studying the effects of varying accretion rates below. Figs. 5, 6, 8 & 11 plot radial profiles of dynamical, magnetic, velocity, and radiation profiles (averaged within annuli, but at a given time). We choose the final time to which our simulation was run for properties at a given time, but highlight how these do (or do not) evolve below. Fig. 14 shows the vertical profiles of different quantities at a given radius. Figs. 7, 9, 10, & 12 show corresponding maps of the field of torques, magnetic fields, velocity fields, and radiation flux fields, while Figs. 4 & 13 show colormaps of other quantities on a wide dynamic range of scales, face-on and edge-on. Figs. 16, 17 compare quantitative profiles versus \dot{m} while Figs. 18, 20, 21 compare projected maps of fields and different scalar quantities versus \dot{m} .

3.1. Convergence to Local Steady-State

Initially in our simulation, there is (by construction) a “cavity” at $R \lesssim 300 R_g$. Gas quickly begins to fill this in, with some inflow through each radius on its dynamical time, but for some time this still leads (as expected) to a surface density $\Sigma_{\text{gas}}(R)$ and inflow rate $\dot{M}(R)$ which decrease with smaller $R \ll 300 R_g$. This means gas mass is building up, and after a time of several order the inflow/accretion timescale in Fig. 5, we see that the central gas profile reaches quasi-steady-state behavior, with approximately constant $\dot{M}(R)$ as a function of radius, out to radii $\gtrsim 10^4 R_g$ where the evolved time since refinement is more comparable to the accretion timescale. This means that if we continued to evolve our simulation even further in time, it is likely that the central accretion rate and gas surface densities would continue to rise until “catching up to” their values at larger radii. But those values are themselves evolving on even longer timescales (Paper I), as the massive gas clouds complex which is being tidally disrupted and fueling the accretion episode continues to accrete, driving evolution in disk properties (like mass and angular momentum direction) on timescales as large as $\sim 10^7 - 10^8$ yr.

We also see that sometimes episodes of strong bipolar outflows, or motions owing to spiral structure, occasionally produce net outflow at some ranges of radii. So it is unclear if the system can ever attain a true “steady-state.” If such a structure is present at radii $\sim 1000 R_g$ at a time when we have refined to $\sim 3 R_g$, and persists for tens of orbits at that radius, it would require evolving our simulations at the ISCO for a time $\gg 10^7 GM_{\text{BH}}/c^3$, which is computationally infeasible. And in this time, other structures with non-steady-state inflow/outflow will appear at other radii. It appears inevitable that at radii where we resolve many orbital times (e.g. the lower-resolution, less-refined simulations from Paper II at $R \gtrsim 300 R_g$, or the simulations here at $\ll 300 R_g$), the accretion rate can systematically increase and decrease by factors of several (and

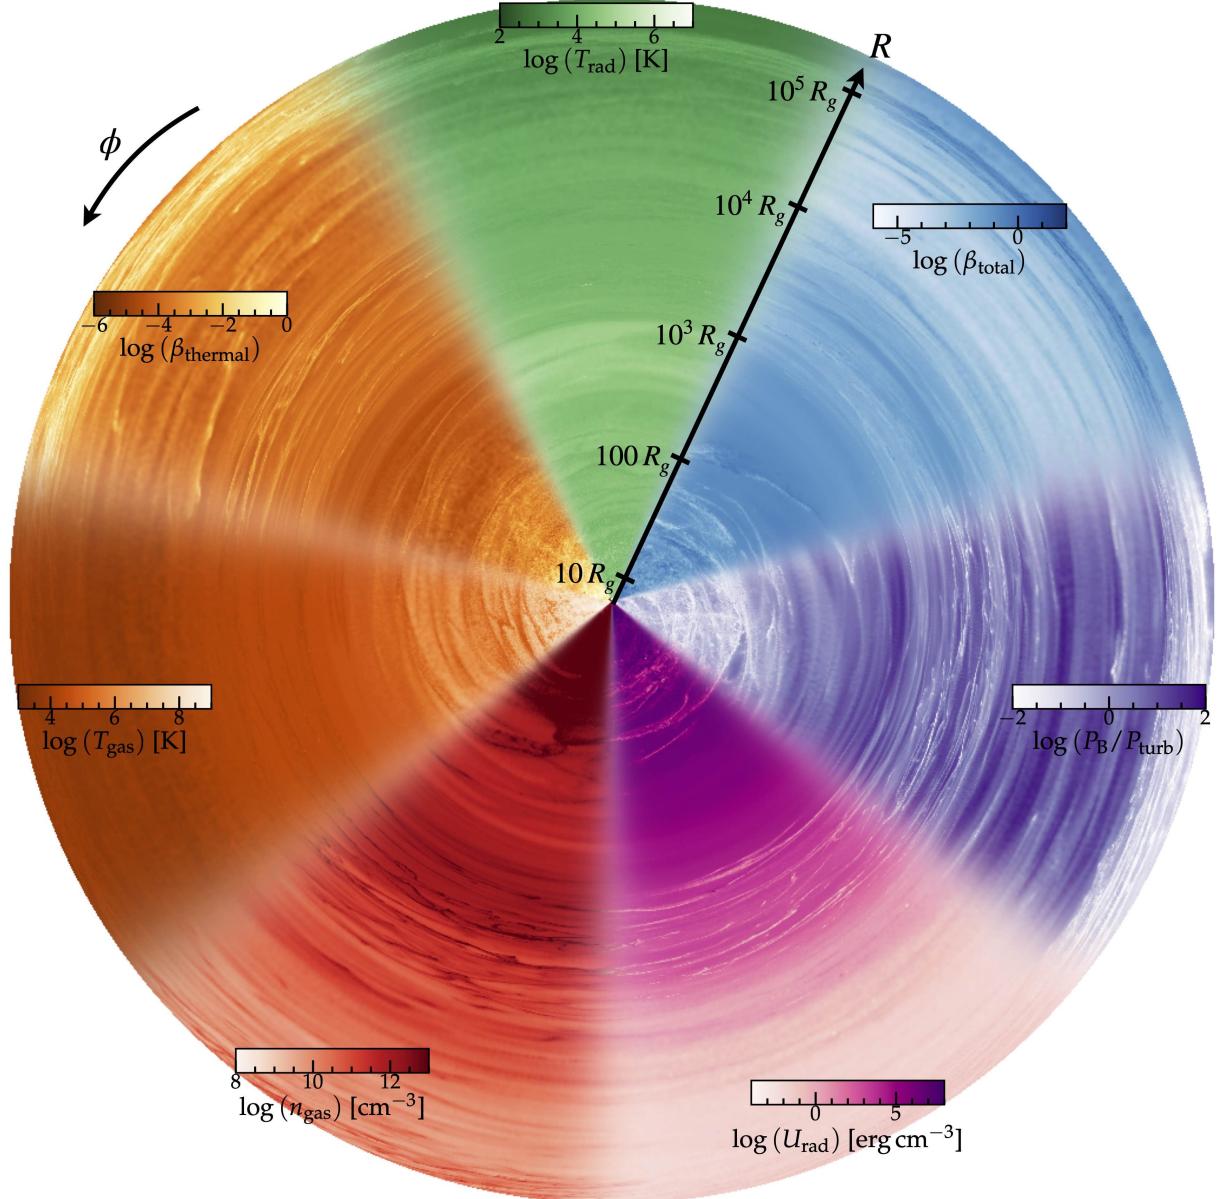


FIG. 4.— Face-on slice through the midplane showing various quantities in the disk in different colormaps. The radial coordinate R has a logarithmic stretch as labeled, from $\sim 3 R_g$ to ~ 1 pc. Different colormaps show: (1) thermal $\beta_{\text{thermal}} \equiv P_{\text{therm}}/P_{\text{mag}}$ (ratio of thermal-to-magnetic pressure), (2) radiation temperature T_{rad} , (3) total $\beta_{\text{tot}} \equiv (P_{\text{therm}} + P_{\text{rad}})/P_{\text{mag}}$ pressure ratio (assuming strong radiation coupling), (4) magnetic-to-turbulent pressure ratio P_B/P_{turb} ($P_{\text{turb}} \equiv (1/2)\rho\delta v_{\text{turb}}^2$), (5) radiation energy density U_{rad} , (6) gas density $n_{\text{gas}} \equiv \rho/m_p$, (7) gas kinetic temperature T_{gas} . As implied by Fig. 1, the disks generally have $\beta_{\text{thermal}} \ll 1$, with $\beta_{\text{tot}} \ll 1$ in the outer disk but more modest $\beta_{\text{tot}} \sim 0.1$ in the innermost disk, and power-law-like gradients in the mean T_{gas} , n_{gas} , U_{rad} . Multi-phase structure is visible in most maps, with trans-Alfvénic turbulence (local P_B/P_{turb} fluctuating below/above unity at all radii).

occasionally order-of-magnitude) over many orbital times, as expected in dynamical accretion models.

It is sufficient for our purposes here that the central regions of interest have reached *local* steady-state. For all our subsequent analysis, we focus on this quasi-steady-state regime after the first few tens of orbits at small radii, as the initial “buildup” is necessarily sensitive to how we initialize the refinement and resolution change. We see these local steady-state conditions maintained for hundreds to thousands of dynamical times at different radii.

3.2. Behaviors Continuing from Larger Radii

Several behaviors we see appear to be natural extensions of what was seen at larger radii in Paper I–Paper II:

Accretion persists. We see (Fig. 5) a quasi-steady \dot{M} in both space and time, with accretion rates $\dot{M} \sim 3 - 6, 0.1 - 0.2$, and $0.005 - 0.01 M_{\odot} \text{ yr}^{-1}$ in the $f_{\dot{M}} = 1, 0.03, 0.001$ runs, respectively (Fig. 16), corresponding to Eddington-scaled $\dot{m} \sim 10 - 20, 0.3 - 0.6, 0.02 - 0.04$, given the BH mass of $\sim 1.3 \times 10^7 M_{\odot}$. We will discuss the luminosities this corresponds to below, but it for the default ($f_{\dot{M}} = 1$) simulation corresponds to a BH growth/mass-doubling timescale of $\sim 0.5 - 2 \times 10^6 \text{ yr}$, which is much shorter than the global duration of the accretion episode expected from the large-scale cloud-disruption dynamics (and seen in the low-resolution

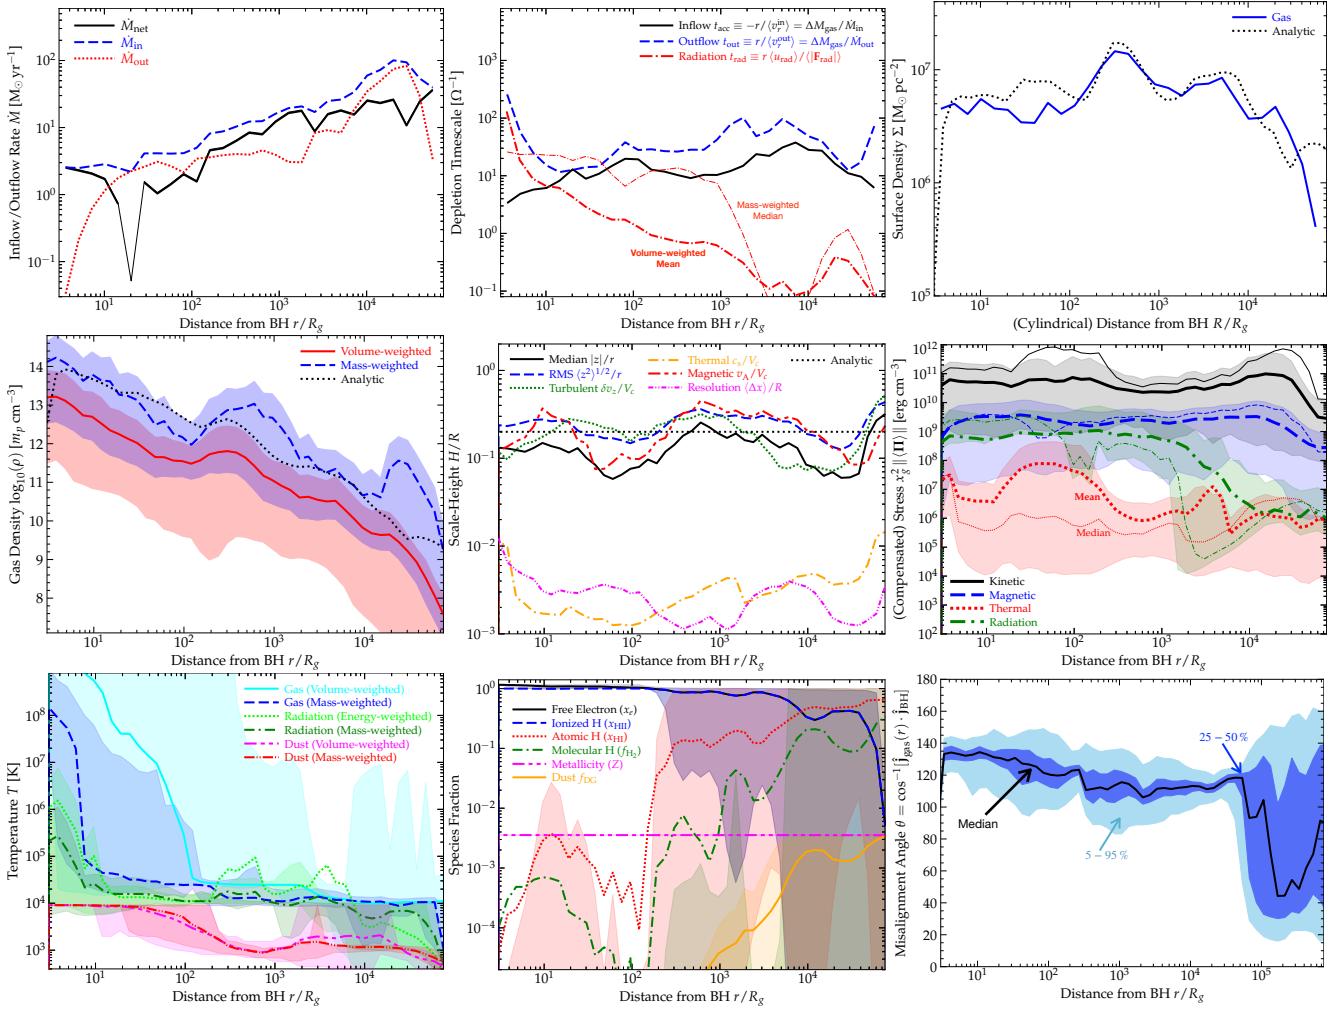


FIG. 5.— Radial profiles of quantities with distance from the BH (versus cylindrical R or spherical r), at the final time in the simulation. (1) Inflow/outflow rate (summing all elements with $v_r < 0$ or > 0) and net inflow rate \dot{M}_{net} through each annulus. Accretion remains highly supercritical with $\dot{M} \sim 3 - 5 M_{\odot} \text{ yr}^{-1} \sim 10 - 20 L_{\text{Edd}}/0.1 c^2$ down to $\sim 3 R_g$. (2) Accretion/outflow (depletion) timescales $\sim \dot{M}/\dot{M}_{\text{disk}}$, in units of $t_{\text{dyn}} \equiv \Omega^{-1} = \sqrt{r^3/GM_{\text{BH}}}$. We also show the radiation escape timescale from the gas $t_{\text{rad}} \sim r \langle u_{\text{rad}} \rangle / \langle |F_{\text{rad}}| \rangle$, with both the volume-weighted mean (dominated by polar escape in low-density channels) and mass-weighted median (approximate escape times from the disk midplane). Accretion is dynamical ($t_{\text{acc}} \sim 1 - 2 t_{\text{orbital}}$), and radiation transport advective in the inner disk. (3) Mean gas surface density $\Sigma_{\text{gas}} \equiv dM_{\text{gas}} / (2\pi R dR)$ in cylindrical annuli, projected face-on. We compare the prediction from the flux-frozen disk similarity model in Paper III (§ 4.4), giving the measured \dot{M} . The agreement suggests the disk similarity is valid. (4) Gas 3D densities $n \equiv \rho/m_p$, volume-weighted and mass-weighted median (lines) and 90% inclusion interval (shaded) shown. We again compare the flux-frozen disk analytic model. Structure arises from spiral arms (Fig. 2), with highly supersonic turbulence generating a broad range of ρ . (5) Disk scale-height H/R (mass-weighted median $|z|$ or rms $\langle z^2 \rangle^{1/2}$), and predicted mass-weighted scale-height from turbulence (vertical $\delta v_z/V_c$), magnetic ($v_A, \phi/V_c$), and gas thermal support (c_s/V_c). We also compare the spatial resolution (Fig. 3), and analytic prediction $H/R \approx v_{\text{turb}}/V_c \approx v_A/V_c \approx \psi \approx 0.2$. The disk is supported by magnetic+turbulent pressure and the scale-heights are well-resolved. (6) Kinetic ($\langle \rho |v|^2/2 \rangle$), magnetic ($\langle |B|^2/4\pi \rangle$), thermal ($\langle \rho u_{\text{th}} \rangle \approx (3n k_B T_{\text{gas}}/2)$), and radiation ($\langle U_{\text{rad}} \rangle$) energy densities (90% range shaded, volume-weighted mean and mass-weighted median in thick and thin lines respectively). Radiation pressure becomes more prominent at small R , but saturates at $\sim 10\%$ of the magnetic pressure (comparable to the turbulent pressure; § 4.1). (7) Gas T_{gas} , radiation T_{rad} , and dust T_{dust} temperatures, volume and mass-weighted (median and 90% range). Reprocessing maintains a relatively cool radiation temperature. (8) Abundances of free electrons (x_e), ionized (HII), atomic (HI), and molecular (H_2) gas, and metal (Z) and dust (f_{DG}) mass-fraction. Dust is mostly sublimated at $T_{\text{dust}} \gtrsim 1400 \text{ K}$ (note T_{dust} is plotted even when dust is sublimated, but has no meaning/impact), but the disk remains highly multi-phase until $\lesssim 100 R_g$. (9) Angular momentum vector and distribution (relative to \hat{j}_{BH} , the pre-existing estimated spin direction of the BH) of the gas within each annulus. Outside the circularization radius, we see a very broad distribution and different mean, as expected, while interior there are still strong warps at the $\sim 10 - 40^\circ$ level. Lens-Thirring precession could strongly modify this in GR.

“parent” simulation from which this particular inflow time was chosen for our zoom-in study). The local gas accretion time is largely independent of accretion rate at $\sim 3 - 20 \Omega^{-1}$ (i.e. a couple orbital times) through most of the disk, though it rises in these units at smaller radii, meaning that the accretion is fundamentally dynamical, rather than secular. In the general state the accretion is clearly not “arrested” (it continues, and is disk-like, per Fig. 1, rather than quasi-spherical or dominated by clumps falling through field lines, though see § 3.5). Instead it proceeds through the midplane (Fig. 9), with $\langle v_r \rangle$

(Fig. 8) as expected given the density profile and accretion rate (Fig. 5).

There is a Thick, Flared Accretion Disk. The disk has a large covering factor/scale height, $H/R \sim 0.1 - 0.5$ (Figs. 2 & 5), weakly dependent on distance from the SMBH and \dot{m} (Figs. 13, 21). But nonetheless it is clearly a disk: the gas is primarily supported in the radial direction by centrifugal motion, on primarily close-to-circular/tangential Keplerian orbits ($\langle v_\phi \rangle \approx V_c(r)$; Fig. 8) with small-to-modest global/coherent eccentricity (Fig. 1, 4, 9). The vertical support of the thick disk

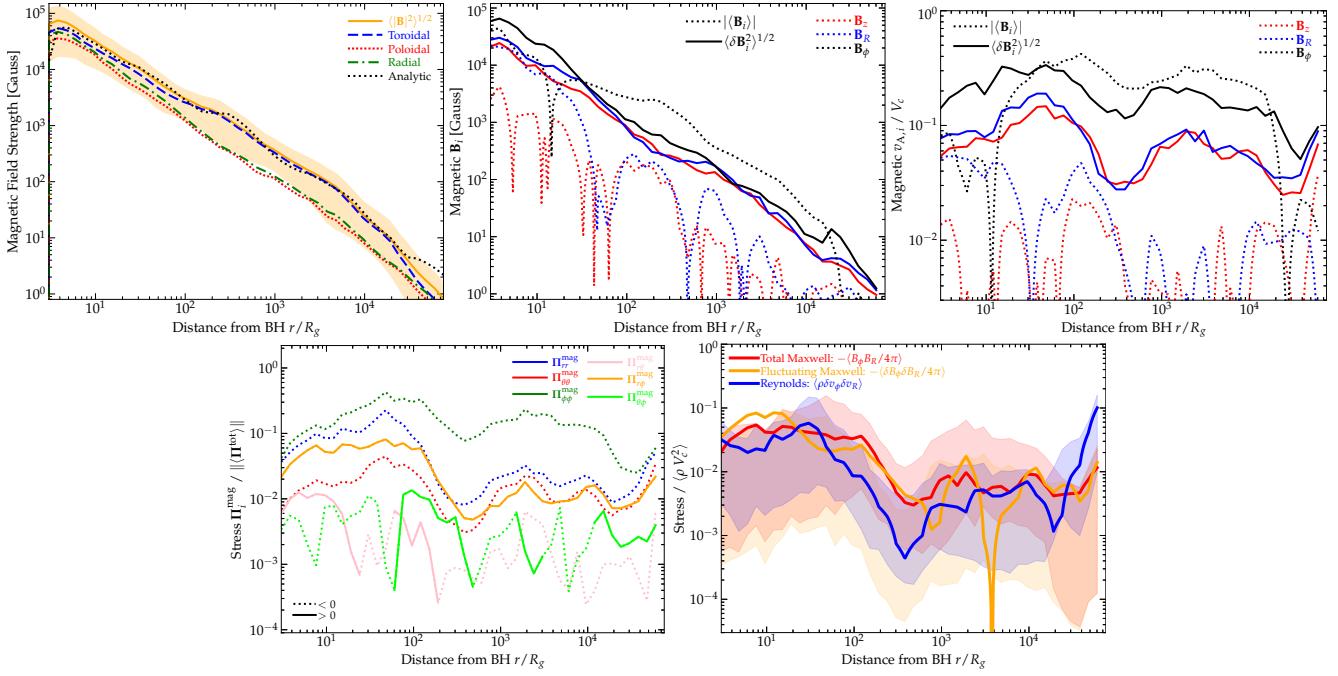


FIG. 6.— Volume-weighted magnetic profiles, as Fig. 5. (1) Magnetic field strengths: volume-weighted rms $\langle B_i^2 \rangle^{1/2}$ toroidal/poloidal/radial ($\langle B_{\phi, z, R} \rangle^2 \rangle^{1/2}$) components and total $|\mathbf{B}|$, with 90% range for $|\mathbf{B}|$. We compare the analytic prediction for $|\mathbf{B}|$ from the same similarity models as Fig. 5, which agree well again. (2) Mean $\langle B_i \rangle$ and standard deviation $\langle \delta B_i^2 \rangle^{1/2}$ of the toroidal B_{ϕ} , poloidal B_z , and radial B_z magnetic fields. The mean toroidal field dominates except in the innermost radii where it is more comparable to fluctuating fields. (3) Mean and dispersion of the Alfvén speeds $\mathbf{v}_{A, i} \equiv \mathbf{B}_i / \sqrt{4\pi\rho}$, relative to V_c . (4) Different components of the Maxwell stress tensor $\mathbf{\Pi}^{\text{mag}} \equiv -\mathbf{B}\mathbf{B}/4\pi + \mathbf{I}|\mathbf{B}|^2/8\pi$, with positive (solid) and negative (dotted) values indicated, normalized to the norm of the total stress tensor from all sources. The toroidal field dominates the magnetic pressure but contributions from fluctuating fields are non-negligible. Angular momentum transport is dominated by the $r\phi$ component, with magnitude comparable to these. (5) Mean and 90% range of the standard total and fluctuating Maxwell stress and Reynolds stress (all comparable), consistent with the accretion rate through each annulus (Fig. 5).

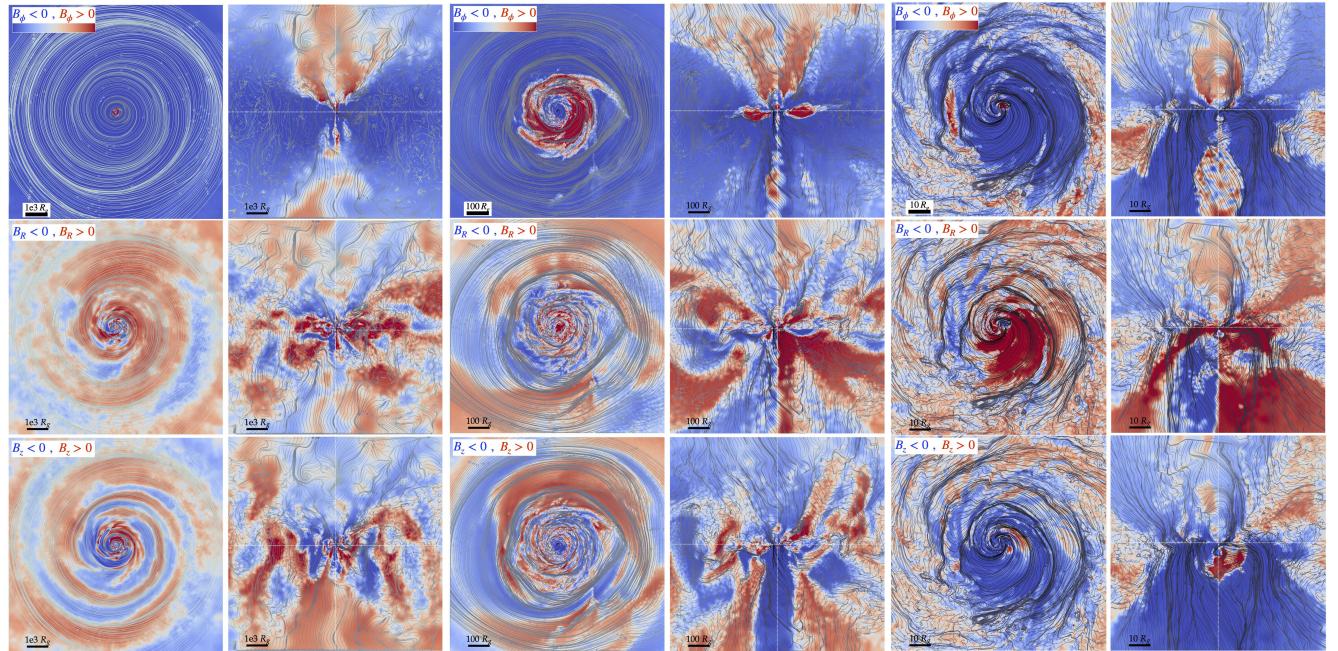


FIG. 7.— Face-on (xy midplane slice) and edge-on (cylindrical Rz wedge) maps of the magnetic field lines in the disk. We show three different scales, from large scales in the disk (left) to intermediate (center) to inner disk (right), as labeled. In each panel, lines show the projected field lines in the plane, while colors encode the sign of the field components labeled: azimuthal B_{ϕ} (top), radial B_R (middle), vertical B_z (bottom). There is a strong mean toroidal field on all scales. A $B_R - B_{\phi}$ anti-correlation continues to all scales, expected from mean-field amplification by induction with the accretion flow and powering strong Maxwell stresses and accretion (Fig. 5). At small radii coherent poloidal fields can thread the midplane from flux carried in from larger R , but they clearly fluctuate and the coherence length of the poloidal field is limited to $\sim H$, which partly explains the lack of large-scale coherent escaping magnetocentrifugal outflows (there is no field with coherence length $\gg R$, and such winds are highly mass-loaded, so will “fail” or “fountain” or “recycle” at heights $|z| \sim H$).

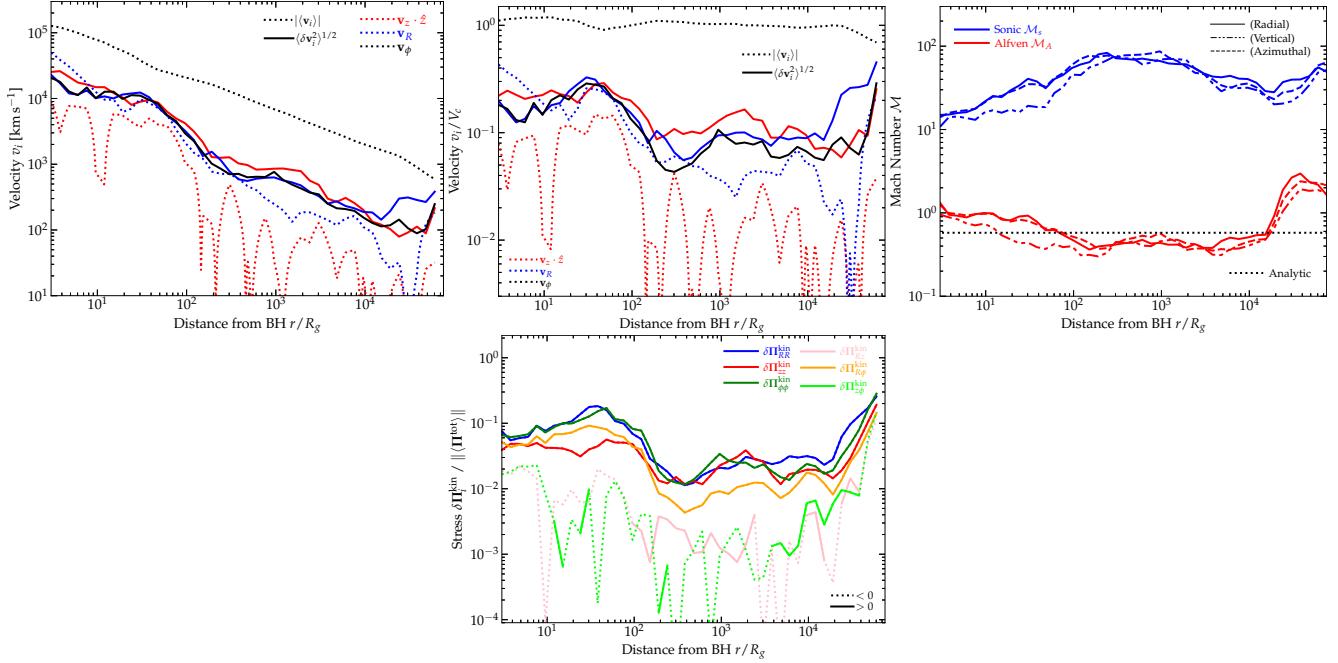


FIG. 8.— Mass-weighted velocity/turbulence profiles, as Fig. 5. (1) Mean $|\langle \mathbf{v}_i \rangle|$ and standard deviation $\langle \delta \mathbf{v}_i^2 \rangle^{1/2}$ of the azimuthal B_ϕ , vertical/outflow $v_z \cdot \hat{z}$, and radial v_R velocity fields. The vertical flow is primarily inflow onto the disk at all radii, while the mean radial advection speed is comparable to the turbulent speed (itself quasi-isotropic) in the inner disk. (2) Same, normalized to the Keplerian velocity V_K at each radius. The thick disks with rapid transport (Fig. 5) follow from the large velocities relative to V_K . (3) Turbulent Alfvén ($\mathcal{M}_A \equiv \delta v/v_A$) and sonic ($\mathcal{M}_s \equiv \delta v/c_s$) Mach numbers, divided into components ($\delta v_{R, \phi, z}$), with the analytic prediction for the 1D $\mathcal{M}_A \approx 1/\sqrt{2}$ from Fig. 5. Turbulence is trans-Alfvénic and highly super-sonic. (4) Different components of the Reynolds stress tensor $\delta \Pi^{\text{kin}} \equiv \rho \delta \mathbf{v} \delta \mathbf{v}$, with positive (solid) and negative (dotted) values indicated, normalized to the norm of the total stress tensor from all sources. The $R\phi$ stress is the dominant source of angular momentum exchange and comparable to the level of the turbulent pressure.

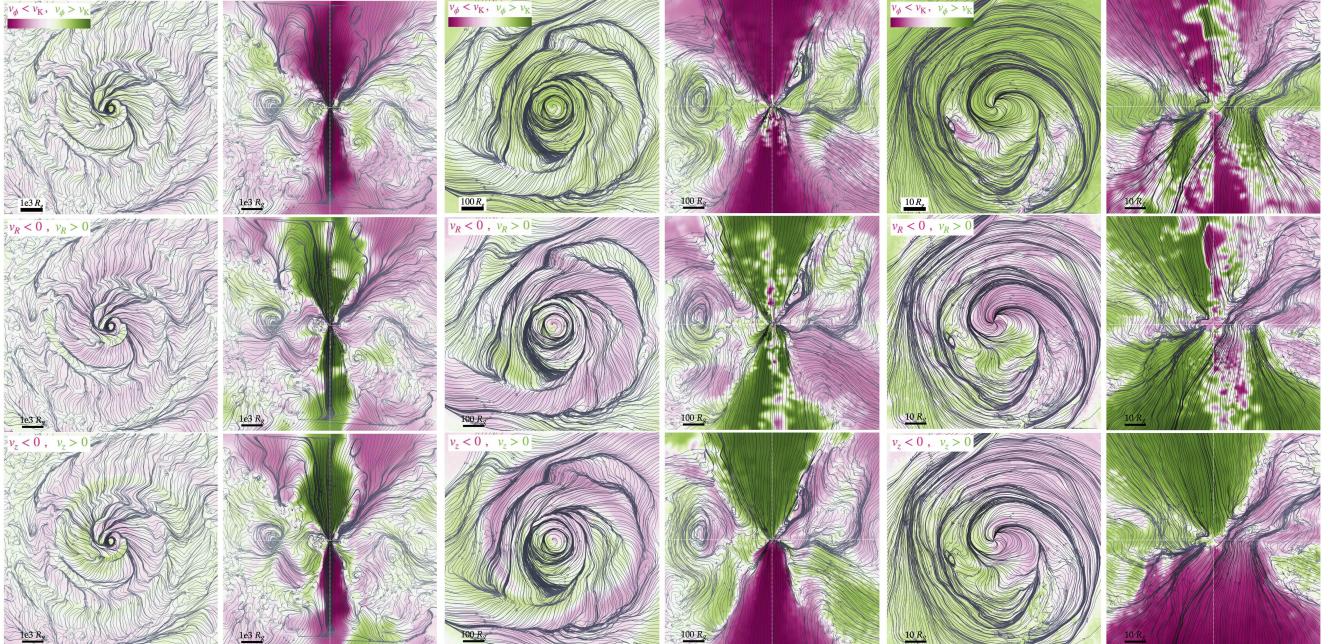


FIG. 9.— Face-on and edge-on maps of the velocity field lines in the disk, as Fig. 7, on different scales. For v_ϕ , we subtract the expected Keplerian velocity for a circular orbit at the given R , $\sqrt{GM_{\text{BH}}/R} \hat{\phi}$, in order to show the more interesting residuals from (primarily) rotational motion. These coincide with the positions of spiral structures in the morphology (Fig. 2). The rotation in the polar region is mostly sub-Keplerian while in the midplane it is mildly super-Keplerian (owing to the Paczyński-Wiita potential and rapid accretion meaning gas does not fully circularize at each radius). Radial and vertical flows are primarily inflow, though both show midplane turbulence and residuals associated with spiral structure.

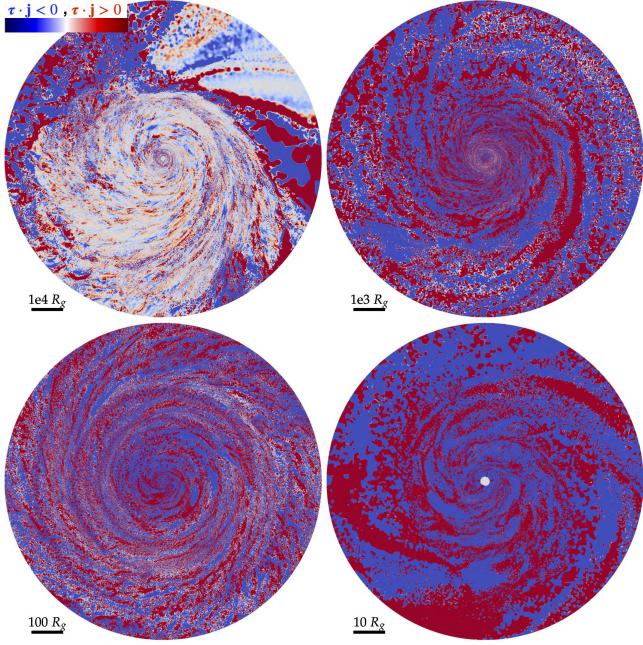


FIG. 10.— Map of the instantaneous torques on each gas parcel, directly from the simulation using the MHD+radiation acceleration computed in that timestep (at the final timestep), projected face-on in an $R\phi$ slide through the midplane on different scales. There is both a small-scale (turbulent) component and a large-scale coherent (spiral) component visible.

comes from a combination of turbulence (with δv_z being trans-Alfvénic and an order-unity fraction ~ 0.3 of $v_{A,\phi}$; Figs. 5 & 8) and magnetic (with gas supported by the in-plane fields $v_{A,\phi}$; Figs. 5 & 6), again all weakly-dependent on \dot{m} (Fig. 16). Interestingly, the systematically smaller median gas-mass-weighted $|z|/R$ appears to closely track the predicted turbulent scale-height $\delta v_z/\Omega$, while the larger rms mass-weighted $\langle z^2 \rangle^{1/2}/R$, reflecting the contribution from the more extended tails of the vertical density distribution visible in Fig. 14 (which also show slow vertical decay of B_ϕ), appears to more closely track the Alfvén scale height $v_{A,\phi}/\Omega$.

The Disk Remains “Hyper-Magnetized.” Specifically the plasma $\beta_{\text{thermal}} = P_{\text{gas, thermal}}/P_B \sim 10^{-5} - 10^{-2} \ll 1$ even in the midplane at all \dot{m} (Figs. 16, 20, 21), as at larger radii (though at the largest \dot{m} , radiation pressure can become more comparable to magnetic, as shown below). The values (see Fig. 5 & 6 for the dispersion in P_B/P_{thermal}) range widely depending on the multi-phase structure of the gas, more driven by large variations in the gas temperature in this structure (hence P_{thermal}) rather than magnetic pressure (though that varies significantly as well, especially as the inner disk thermalizes and becomes warmer and hotter coronal gas becomes more prevalent above/below the disk at smaller radii). As we show below, removing magnetic fields from the inner disk (even while retaining radiation pressure) produces qualitatively different outcomes and disk structure.

Outflows are Modest. We see (usually radiation pressure-driven) outflows which are not negligible at some radii, but are in general modest, usually comprising tens of percent to an order one fraction of the inflow rate (though see § 4.7). As discussed further there are transient phases where the outflow rate at some radii exceeds the local inflow rate, though they do not generally exceed the inflow rate from the largest radii $\gtrsim 10^5 R_g$ (around the BHROI) where the inflow rates are as large as $\sim 30 M_\odot \text{ yr}^{-1}$. So they typically do not change

the mass or accretion structure of the disks themselves at the orders-of-magnitude level, though they are certainly of interest for AGN “feedback.” Note that *some* of the apparently large outflow rates in Fig. 5 (e.g. the apparent large outflow “spike” at $\sim 3 \times 10^4 R_g$) can clearly be seen in Figs. 8, 9, 10 to *not* actually be large-scale winds, but rather related to in-plane disk motions such as spiral/eccentric modes (plainly visible in Fig. 1) in the midplane causing local inflow-outflow sloshing/epicyclic motions (e.g. non-zero oscillating v_r), while the mean sense of the radial velocities is largely inflow, especially on large scales. The true “outflows” in the sense usually meant will be explored more below but tend to be bipolar and composed of more tenuous gas. This is broadly consistent with previous simulations even of highly super-Eddington, radiation-pressure-dominated inner accretion disks, which generally see $\dot{M}_{\text{out}} \sim \dot{M}_{\text{in}}$ (Jiang et al. 2019). We stress that we do not include line driving or GR/spin effects, nor evolve the simulations long enough to follow wind-launching from large scales (e.g. the torus or narrow-line region), so the only outflows we can follow are driven by magnetic fields and/or continuum radiation pressure and almost certainly under-estimate the true magnitude of outflows, as discussed further below (§ 4.7).

The Disk Features Radiative, Trans-Alfvénic Turbulence. We see that both velocity (Fig. 8) and magnetic field fluctuations (Fig. 6) are roughly isotropic with Alfvén Mach number of order unity, remarkably constant as a function of radius and \dot{m} (Fig. 16). Midplane fluctuations with coherence lengths $\sim H$ are clearly visible in all cases (Figs. 4, 7, 9, 13, 18). Since $\beta_{\text{thermal}} \ll 1$, this necessarily means the sonic Mach number $\mathcal{M}_s = v_{\text{turb}}/c_s \gg 1$, so the turbulence is highly supersonic, enabled by efficient cooling ($t_{\text{cool}} \ll \Omega^{-1}$), leading to compressive shocks which feature large density fluctuations (up to $\sim \mathcal{M}_s^2$, producing the variance seen in midplane density at a given radius) and multi-phase structure (density and temperature spreads in Fig. 5) and visible as well in the extremely small-scale density structure in Figs. 1-2 reaching down to scales of order the expected sonic length/shock width/Sobolev length (all dimensionally similar at size $\sim H/\mathcal{M}_s^2 \sim c_s^2/v_A \Omega$). As shown in § 4.9, the scattering mean-free-paths of photons even in the inner disk at the highest \dot{m} here are generally larger than the intrinsic shock width, so radiation pressure does not strongly alter our intuition.

Accretion Is Driven by In-Plane Maxwell and Reynolds Stresses. The torques on the gas are clearly turbulent/local on scales $\sim H$ (Fig. 10), and the Maxwell and kinetic/Reynolds stresses which involve angular momentum transport are dominated by the usual in-plane $R\phi$ terms (Figs. 6 & 8) for all \dot{m} (Fig. 17). As in Paper II, the net accretion rates \dot{M} , inflow speeds $\langle v_R \rangle$, and torques $\langle \tau \rangle$ are all consistent with the magnitude of the Maxwell ($-\langle B_\phi B_R/4\pi \rangle$) and Reynolds ($\langle \rho \delta v_\phi \delta v_R \rangle$) stresses, with the Maxwell stresses generally larger by a factor of a few in the inner disk compared to the Reynolds stresses and, notably, the *mean* Maxwell stress ($-\langle B_\phi \rangle \langle B_R \rangle / 4\pi$) comparable to or larger than the fluctuating Maxwell stress ($-\langle \delta B_\phi \delta B_R \rangle / 4\pi$). The similarity of these across values of $f_{\dot{M}}$ (Fig. 17) suggests that the basic similarity arguments in Paper III provide a reasonable first-order description of the disk properties.

Spiral Structure Persists. There are large-amplitude, clearly visible spiral ($m = 1$) modes throughout in all cases, as noted above (Figs. 1, 4, 7, 9, 10). These evolve in time but are always visually present and appear at all radii in the disk. Decomposing the disk into non-axisymmetric modes,

the dominant mode is indeed $m = 1$, as expected for a gravitationally stable (non-fragmenting, negligibly self-gravitating) disk in a close-to-Keplerian potential (Tremaine 2001; Jacobs & Sellwood 2001; Touma 2002; Bacon et al. 2001; Hopkins & Quataert 2011b; Hopkins 2010; Hopkins & Quataert 2011a). And indeed these have also been seen in many other accretion disk simulations on these scales (Heinemann & Papaloizou 2009; Hopkins & Quataert 2010b,a; Hopkins et al. 2016; Gaburov et al. 2012), especially of high-accretion rate or super-Eddington disks even without any self-gravity (Jiang et al. 2019; Kudoh et al. 2020; Davis & Tchekhovskoy 2020; Kaaz et al. 2022).

Star Formation Is Completely Shut Down. At all \dot{m} and all radii we analyze (\ll pc scales), the disk has Toomre $Q \gg 1$, reaching $\gg 10^6$ if we include magnetic support (not explicitly shown, but easily inferred from the properties plotted in Fig. 5, and plotted at larger radii in Paper II). The magnetic critical mass is also much larger than the entire disk mass. So we see star formation is strongly suppressed relative to inflow at all radii \lesssim pc, and zero stars form over the duration of the simulation at radii $r \lesssim 0.05$ pc. Only a couple stars have orbits taking themselves into radii $\lesssim 0.01$ pc, and these (as well as their winds/jets/radiation) have a negligible effect on the dynamics (see Hopkins et al. 2024a, for details). This is not trivial: if we assumed an SS73-like α disk with $\alpha = 0.01 - 0.1$ and the same BH mass and accretion rate, we would predict $Q \ll 1$ at all radii $\gtrsim 200 - 500 R_g$.

Magnetic Fields Continue to Scale with Radius. We see growth of the magnetic fields in a Lagrangian sense as gas accretes inwards, giving a spatial scaling typical $|\mathbf{B}| \propto \dot{m}^{1/2} r^{-1}$ very approximately (Fig. 5, 16), or $|\mathbf{B}| \propto H^{-1}$ or $\rho^{1/4} r^{-3/4}$ at a given \dot{m} (Fig. 6). This is consistent with the expectations from simple models of \mathbf{B} supplied by flux-freezing or net advection of magnetic flux from larger radii, with a large-scale mean toroidal field (see Paper II), although the actual physics governing the mean (let alone fluctuating/turbulent) field can be much more complicated (see Johansen & Levin 2008; Gaburov et al. 2012; Bai & Stone 2012, plus Squire et al. 2024, Guo et al., Tomar et al., in prep.). Note at larger radii, these scalings are approximately the same as the simpler $B \propto \rho^{2/3}$ adopted in Paper III; here we see the latter is not quite as accurate an approximation at small radii and large \dot{m} where radiation pressure modifies the profiles, which we discuss below. The mean toroidal and radial fields are anti-correlated (Fig. 6 & 7), again as expected.

The Magnetic Field Remains Predominantly Toroidal. The field is ordered, with the clearly-dominant component being the midplane toroidal/azimuthal field at most radii (Fig. 6) and \dot{m} . There is a non-negligible but smaller mean radial field. And there is a quasi-isotropic turbulent/fluctuating field, of comparable amplitude to the mean radial field, $\sim 10\%$ of the mean azimuthal field (Figs. 6, 7). The mean/coherent vertical component through the midplane generally remains the weakest component. Again this fits the expectation from Paper II for accretion of gas into the disk initially via streamers or other in-plane structures, which amplify the toroidal-radial mean field and weaken the vertical field. This places the mean vertical/poloidal field in the “turbulence-dominated” regime (so turbulent resistivity suppresses formation of a large coherent mean vertical field). But it is striking that this coherent toroidal field structure persists at all \dot{m} even in the strongly turbulent, radiation-pressure-dominated zone. However, in § 3.5 we show note that the field can change in episodes associated

with strong outflows.

All of these are qualitatively consistent with analytic extrapolations in the simple models Paper III, from the behaviors in Paper I-Paper II from larger scales.

3.3. Behaviors That Change In the Inner Disk

Unlike Paper I-Paper II, we now resolve the transition from a disk with negligible to important radiation fields, in all three of our different accretion-rate cases.

The Disk Self-Ionizes. As temperatures rise towards the inner disk and the local radiation field becomes harder, the disk transitions from primarily atomic to fully ionized. The ionized gas mass fraction begins to rise significantly, but there is still a mix of ionized and atomic hydrogen, interior to $\lesssim 2 \times 10^{16}$ cm ($\lesssim 10$ light-days or $\lesssim 6000 R_g$). This rises steadily until at $\lesssim 10^{15}$ cm or ~ 0.5 light-days ($\lesssim 300 R_g$) the atomic mass fraction drops steeply, in our default ($f_M = 1$) simulation, though the radius where this occurs appears to scale with accretion rate (Fig. 16). For each of our different \dot{M} runs, these radii are within a factor of ~ 2 of the analytically-predicted radius for the same M_{BH} and \dot{M} for the simplified magnetically-dominated disk models in Hopkins (2024).

The Flow Becomes Increasingly Radial. The inner regions at $\ll 300 R_g$ exhibit an increasing $\langle v_R \rangle$ and radial Reynolds stress Π_{RR} (Fig. 8) which also manifests in stretching the magnetic field lines to enhance B_R (Fig. 6, 17).

Radiation Pressure Saturates at Similar Levels to Turbulent and Magnetic Pressures. At radii $\gtrsim 0.01$ pc $\sim 10^4 R_g$, $P_{\text{rad}} \sim P_{\text{therm}} \ll P_{\text{turb}} \sim P_{\text{mag}}$ – i.e. radiation pressure is order-of-magnitude similar to thermal pressure and both are negligible dynamically. Relative to other pressures/energy densities in the disk, we see P_{rad} and U_{rad} rise as expected (Figs. 4, 5, 13), so $\beta_{\text{tot}} \equiv (P_{\text{rad}} + P_{\text{therm}})/P_B$ increases,⁵ until saturating at values $P_{\text{rad}} \sim 0.1 - 1 P_B$ or so, i.e. radiation energy density comparable to the turbulent or fluctuating magnetic energy density, but notably a factor of a few smaller than the mean magnetic energy density. There is substantial variation in P_{rad} in the multi-phase regions of the disk and between the midplane and polar regions (Figs. 4, 13). Because the simulations with different \dot{m} have radiation pressure roughly $\propto \dot{m}$, and magnetic pressure $\propto \rho v_A^2 \propto \dot{m}$, this ratio does not depend strongly on \dot{m} (Fig. 16, 20, 21).

Energy Transport Becomes Advective at the smallest radii. Per Fig. 11, in the outer portions of the disk, $R \gtrsim 10^{-3}$ pc $\sim 3 \times 10^{15}$ cm $\sim 1d \sim 800 R_g$, the radiation flux escaping from the disk is the primary energy loss mechanism, and corresponds to the expected gravitational energy flux in a Keplerian thin disk, $F_{\text{rad}} \sim F_{\text{grav}} \sim (3/8\pi)\dot{M}\Omega^2(1 - \sqrt{3R_g/R})$ (Shakura & Sunyaev 1973). This is comparable to but somewhat larger than the transport of energy through the disk, an order-unity fraction of which goes into compression of the toroidal field, since $B_\phi^2 \propto R^{-(2-3)}$ rises rapidly towards smaller radii; see Paper III and Figs. 5. And the radiation energy densities in the midplane are higher than F/c by a factor of $\sim \tau$, more or less as expected (Figs. 11, 16). Note

⁵ Here and throughout, we define $P_{\text{rad}} \equiv u_{\text{rad}}/3$, and $\beta_{\text{tot}} \equiv (P_{\text{rad}} + P_{\text{therm}})/P_B$, which follow standard literature conventions. However we caution that the interpretation of β_{tot} and P_{rad} depends on the degree of radiation-gas coupling: P_{rad} will only act like a “pressure” dynamically in the local tight-coupling (infinitely small radiation mean-free-path) limit, which we show in § 4.9 is rarely true in the simulations. A more meaningful comparison of the actual dynamical importance of radiation and magnetic fields or gravity is manifest in the salient forces/accelerations, discussed in § 4.

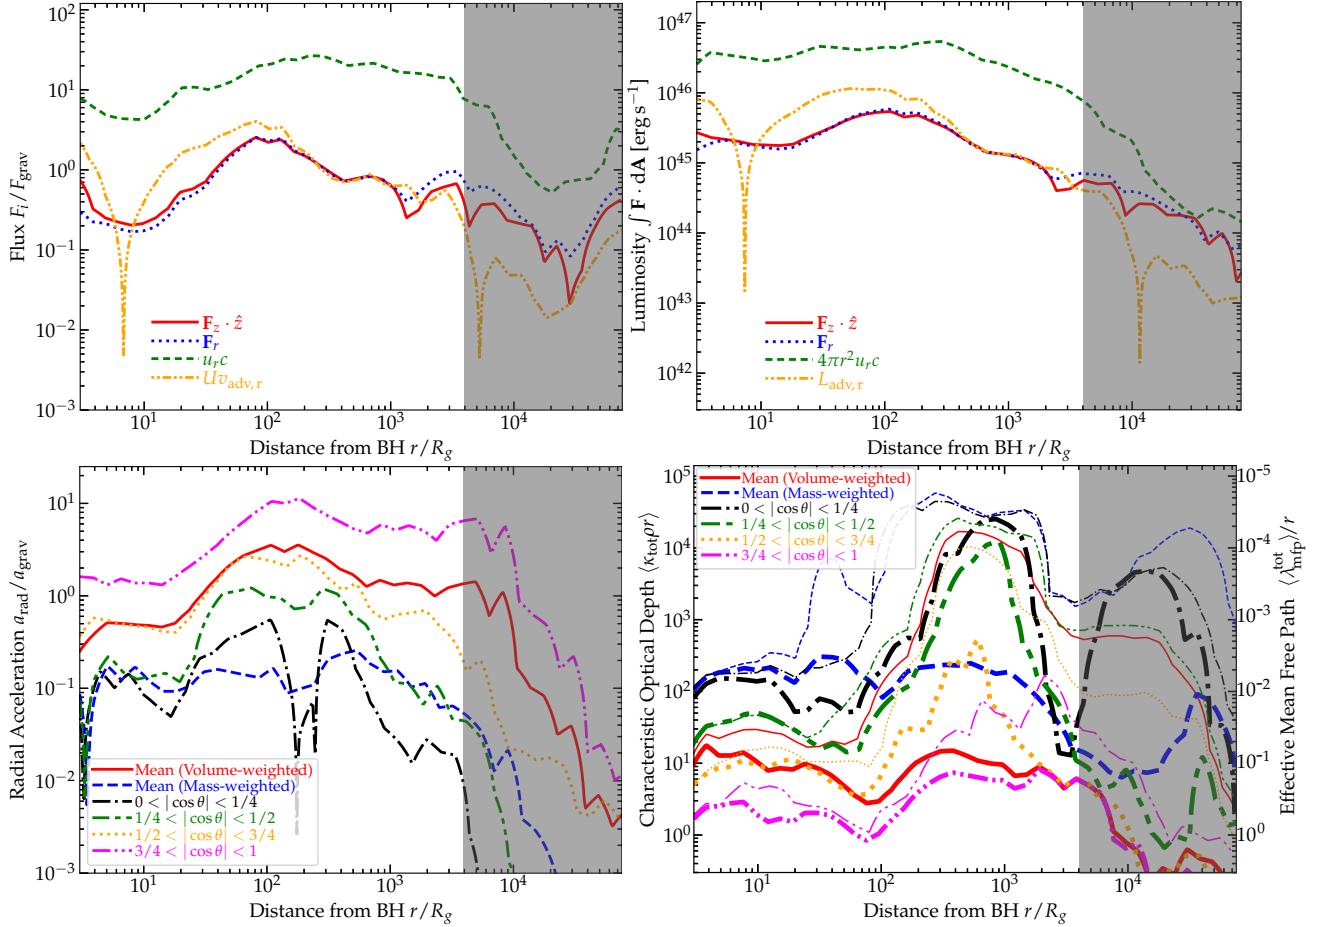


FIG. 11.— Volume-weighted radiation profiles, as Fig. 5, for our high- \dot{M} ($f_{\dot{M}} = 1$) simulation. (1) Mean flux \mathbf{F}_i by vertical/radial components, normalized to the total flux expected for a radiatively efficient thin disk truncated at the ISCO: $F_{\text{grav}}^{\alpha} \equiv 3/(4\pi)\dot{M}\Omega^2(1 - \sqrt{3R_g/R})$. We also compare the “optically thin” flux/radiation energy density $u_{\text{rad}} c$, and the advective flux of total internal energy $v_r (e_{\text{rad}} + e_{\text{therm}} + e_B)$. The disk becomes radiatively inefficient and advection-dominated at small radii. (2) Same, but the integral of each in thin shells to represent the total energy flow rate per unit time. The standard Eddington luminosity $L_{\text{Edd}} = 3.2 \times 10^4 L_{\odot} (M_{\text{BH}}/M_{\odot}) = 1.6 \times 10^{45} \text{ erg s}^{-1}$ here. The luminosity saturates modestly super-Eddington (though we do not resolve the innermost near-horizon regime). (3) Radial acceleration $|\mathbf{a}_{\text{rad}} \cdot \hat{\mathbf{r}}|$ relative to gravity $|\mathbf{a}_{\text{grav}} \cdot \hat{\mathbf{r}}|$ (computed in-code), weighted by gas volume or mass or volume-averaged in different wedges of polar angle $|\cos \theta| (= 1 \text{ along the angular momentum axis})$. In the polar direction radiation exceeds gravity, especially at somewhat larger radii $\gtrsim 100 R_g$, but in the midplane (where most of the gas mass resides) radiation pressure is weak. (4) Characteristic optical depth or photon mean-free-path $\langle \kappa_{\text{tot}} \rho r \rangle$, estimated as $\langle \kappa_{\text{tot}} \rho r \rangle \sim r / \langle \lambda_{\text{mfp}}^{\text{tot}} \rangle$ as a function of polar angle and radius (in terms of total absorption+scattering opacity). We show both mass-weighted (or volume-weighted where noted) median opacities (thin), typical of the disk midplane, and Rosseland mean opacities in each shell (thick), more representative of what is seen by escaping radiation. Shaded radii show where the radiation equilibration timescales $\sim \tau_{\text{tot}} r / c$ tend to exceed ($\gg 10^5 GM_{\text{BH}}/c^3$) our total runtime since reaching maximum refinement near the ISCO (where much of the radiation is produced).

that most of the total flux (Fig. 11) at large radii is the flux from the inner disk, which both illuminates the outer disk and represents escaping radiation in the polar regions. But interior to $R \ll 1 \text{ ld}$ in our supercritical (highest accretion rate, $f_{\dot{M}} = 1$) simulation, and within $\lesssim 10 R_g$ in the sub-critical simulations (approaching the ISCO), we see the energy transport becomes primarily advective, and the escaping radiation flux falls increasingly below F_{grav} (Fig. 16). We also see this in the accretion timescales becoming comparable to or shorter than radiation escape timescales in Figs. 5, 16. Note, per Fig. 5, that most of the advected energy density is not in radiation: summing the energies of advected turbulent motions, radiation, and magnetic fields, this is largely dominated by the advected magnetic mean-field (B_{ϕ}) energies (Fig. 5), at these radii.

Radiation Escapes in a Widening Bicone & Illuminates the Disk. Fig. 12 and Figs. 18, 21 show that the radiation flux escapes in a broad-angle bicone, the effective opening

angle of which appears to broaden at larger radii. The disk midplane is shielded, and features very little net flux in the Lagrangian frame (Fig. 11), but the surface layers are strongly illuminated at all \dot{m} (Fig. 18). This owes partly to the thick, flaring disk structure (Figs. 2, 13, 21) but also to indirect scattering onto the disk. We see this in Fig. 12 as flux streamlines at intermediate angles appear to “curve down” before hitting the disk, a behavior which is more dramatic at higher \dot{m} . This occurs because of scattering in the lower, but still significant gas densities above the disk (and those gas densities are higher at higher \dot{m}), where the gas is mostly ionized and so has a high scattering opacity but densities slowly decreasing with $|z|$ (Fig. 14). The effective size of the emitting region is therefore large: we can see from Fig. 12 that a significant ($O(1)$) fraction of the sightlines from the inner disk will have their flux reprocessed at radii as large as the outermost disk ($\gtrsim 10^{16} \text{ cm}$). This is confirmed in Kaaz et al. (2024), and will be modeled in more detail in future work (Bardati et al., in prep.).

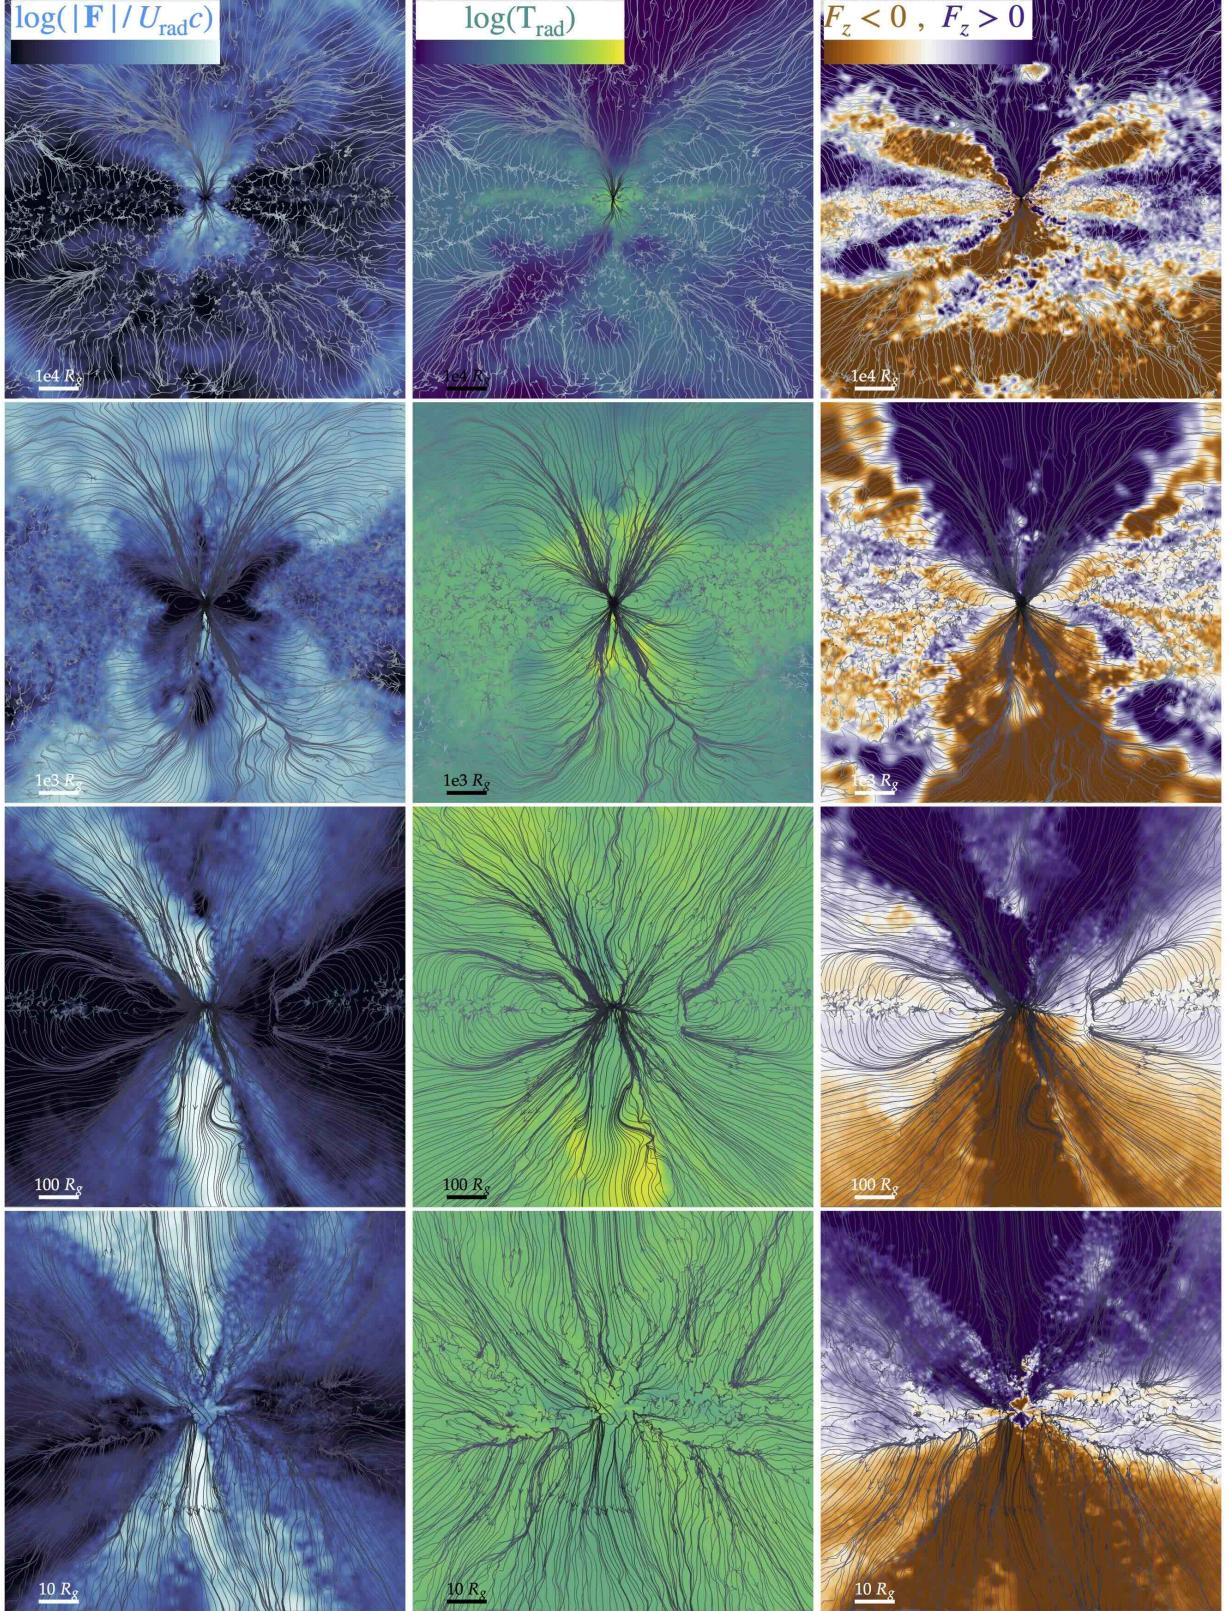


FIG. 12.— Radiation flux maps, as Fig. 7. Lines show radiation streamlines (following the flux vectors \mathbf{F}_{rad}), edge-on (R_z) in a cylindrical wedge, with box side-length decreasing *top-to-bottom*. Fluxes are summed over all bands followed, and defined in the Lagrangian frame (so does not include the advective flux). *Left:* Colors show $|\mathbf{F}_{\text{rad}}|/U_{\text{rad}}c$, from 0 (trapped; black) to 1 (free-streaming; white) on a logarithmic scale. *Center:* Colors show radiation temperature, as Fig. 4. *Right:* Colors show sign(F_z), as Figs. 7 & 9. In the polar direction radiation escapes, but along the disk surface much of the radiation scatters into or is absorbed into the disk, which is externally heated/irradiated at $R \gtrsim 100 R_g$ so the outer disk flux vectors flow *into* the disk.

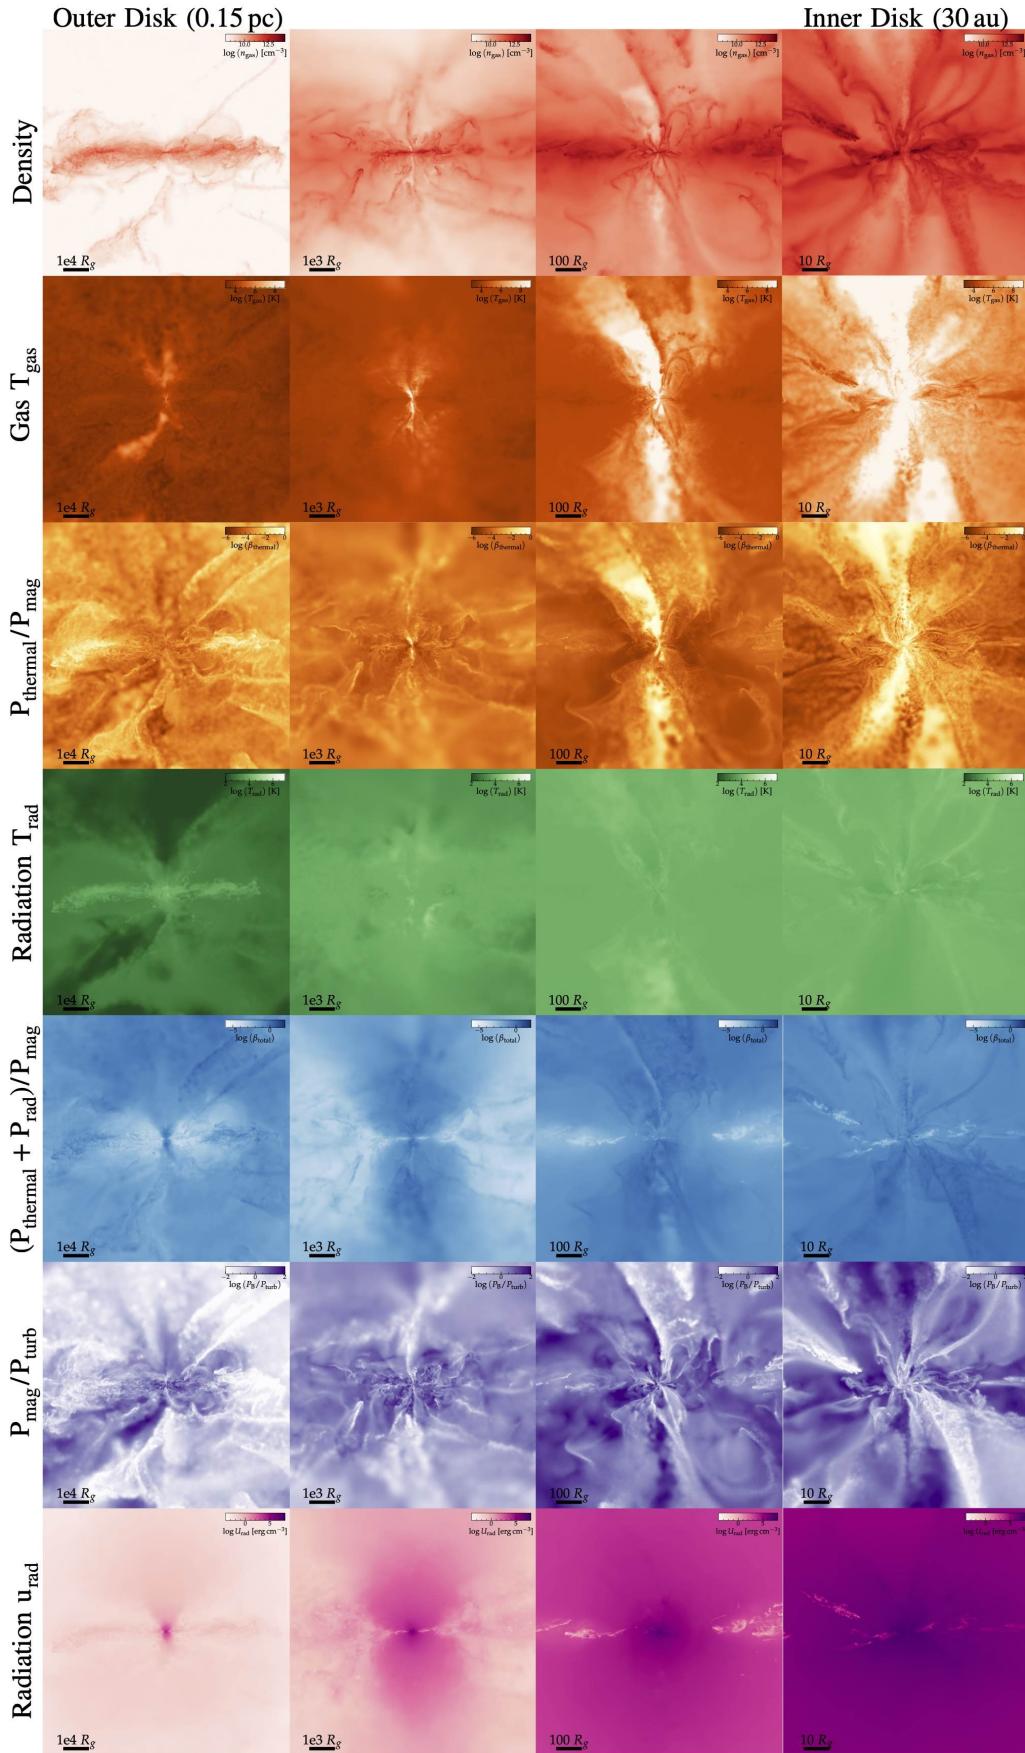


FIG. 13.—Edge-on R_Z cylindrical wedge maps of quantities from Fig. 4, in different box sizes (zooming in left-to-right).

There Is An Extended Scattering “Photosphere.” We see in Figs. 11 & 12 that there is significant scattering, with deflection of photon trajectories and slowing the escape of photons (making $u_{\text{rad}} \gg |\mathbf{F}_{\text{rad}}|/c$) even in the broad regions above the disk in the innermost regions. Measuring the mean column densities in different polar angles θ , we see that the scattering optical depth (τ is primarily from electron-scattering here) in our highest- \dot{M} simulation remains $\sim 5 - 10$ even along a broadly polar set of angles $3/4 < |\cos \theta| < 1$, out to a thousands of R_g , i.e. $\sim 0.5 - 1 \times 10^{16} \text{ cm}^{-3}$. This helps explain the scattering “down” into the disk described above, and enhances the reprocessing of radiation by the outer disk as well as providing a sort of “warm absorber”-like Comptonizing medium for the continuum radiation. In our intermediate $f_{\dot{M}} = 0.03$ simulation with $\dot{m} \sim 0.3 - 0.5$, the optical depths along the polar direction or in a volume-weighted sense are systematically smaller (as expected, owing to the systematically lower gas mass supply and therefore densities at all radii), but they remain non-negligible (order-unity) out to similar radii (Fig. 16). Only in our lowest accretion rate case ($f_{\dot{M}} = 0.001$) with $\dot{m} \sim 0.02$ (where the disk itself is only marginally optically-thick) do we see optical depths $\ll 1$ along the polar/volume-filling directions. Still, even in this case, a non-negligible fraction of the light from the inner disk ($\gtrsim 10\%$) will be either scattered down from polar regions (with optical depths $\sim 0.01 - 0.1$) or intercepted by the outer disk (which still features $H/R \sim 0.1 - 1$ at large radii; Fig. 21).

Emergent Radiation Temperatures are Modest. In the innermost regions, there are regions where the radiation temperature becomes quite large, $\sim 10^6 \text{ K}$ (Figs. 5, 16). However, this is strongly reprocessed and, because of the radiative efficiency being such that most of the emission comes from a very broad range of radii, means that the emergent luminosity (even along the vertical “cone” in Fig. 12), has radiation temperatures which are much more modest (Figs. 4, 5, 13). Indeed, in the direct escaping “cone”, the closest to a “direct” sightline to the inner disk (Fig. 12), the escaping radiation at the furthest distances photons have diffused here gives an emergent effective blackbody temperature of $\sim 1 - 3 \times 10^4 \text{ K}$ (Fig. 5) where it is escaping vertically from the outer thermalized radii of the disk (see Fig. 16, 21 for different \dot{m}). Note that in the outer disk, the effective temperatures appear to drop below $< 10^4 \text{ K}$. However, examination of the optical depths, gas temperatures, and chemical structure of the disk (e.g. Fig. 5) shows that this occurs at radii where the disk becomes primarily atomic, dust is no longer sublimated, and one can no longer assume local blackbody equilibrium $T_{\text{rad}} \neq T_{\text{gas}} \neq T_{\text{dust}}$. Since our effective radiation temperature evolved in-code is determined by separately evolving the photon energy density and number density, it is a *photon number weighted* effective radiation temperature. That means that even a small fraction ($\sim 1 - 10\%$) of the optical/UV radiation being reprocessed by dust and re-emitted with temperatures $T_{\text{dust}} \sim 100 - 1000 \text{ K}$ will drop the weighted radiation temperature to $T_{\text{rad}} \ll 10^4 \text{ K}$. We expect a more detailed frequency-resolved radiation treatment would give rise to a highly multi-component spectrum in this regime (with separate contributions to the emergent QSO spectrum in escaping X-ray/EUV, UV/optical, NIR, and mid-far IR), but modeling this requires more detailed post-processing Monte Carlo radiation transport (Bardati et al., in progress). Given that this is happening embedded in a dusty torus and massive, high-redshift starburst galaxy (with galaxy-integrated star formation rate of $\sim 200 \text{ M}_{\odot} \text{ yr}^{-1}$), it is likely that on large ($\gg \text{pc}$

scales the system will remain isotropically obscured and the radiation will be further downgraded, but this still gives us a useful proxy for the “intrinsic” quasar SED.

Some of these behaviors are anticipated, or at least qualitatively similar to standard thick or “slim” disk models for supercritical accretion (Paczynsky & Wiita 1980; Abramowicz et al. 1988). This is despite the fact that slim-disk models generally neglect magnetic pressure in the dynamics and energy transport of the disks entirely (the fields only appear implicitly in some assumed stress), and instead begin from assumptions like those in standard thin thermal-pressure-dominated α -disks (Shakura & Sunyaev 1973). Comparison to analytic models is discussed further below.

3.4. Luminosities & Radiative Efficiencies

We see in Fig. 16 that for $\dot{m} \lesssim 1$, the bolometric luminosities L_{bol} are dominated by the energy from the central $\sim 10 R_g$,⁶ and the inflow/accretion rates \dot{M} are roughly constant with R within the inner disk (central $\sim 100 - 1000 R_g$). From direct comparison of these, we can estimate the radiative efficiency, $\epsilon_r \equiv L_{\text{bol}}/\dot{M} c^2$, which is $\sim (0.15, 0.1, 0.02)$ for $f_{\dot{M}} = (0.001, 0.03, 1)$ ($\dot{m} \sim 0.02, 0.3, 20$). For $\dot{m} \lesssim 1$ these are similar to the canonical expected efficiency of any radiatively-efficient rotating disk, though we caution that in a full GRMHD calculation these could change at the order-unity level (as we only adopt a simplified PW potential).

For $\dot{m} \gg 1$, a corollary of the point in § 3.3 that the flux becomes advective at radii well outside the ISCO is that the system becomes less radiatively efficient. Fig. 11 shows that the radiation fluxes fall well below F_{grav} at the smallest radii, and correspondingly the emerging luminosity has roughly equal contributions per logarithmic radius R (so the total emergent luminosity grows logarithmically only as gas migrates further in, rather than increasing $\propto R^2 F_{\text{grav}} \propto R^{-1}$). The net luminosity emerging from the thermalized disk is \sim a few $\times 10^{45} \text{ erg s}^{-1}$ (a couple to a few times Eddington), more weakly dependent on the cutoff/innermost radius suggesting (seen visually in how the luminosity rises with radius weakly in Fig. 11). This is consistent with the radiation interior to a trapping radius further out being advected inwards, also consistent with the flow seen in Fig. 12.

3.5. Variability, Bursts, & Magnetic State Changes

All of our simulations exhibit some variability in accretion rates, luminosities, and time-and-space-dependent fluxes and radiation temperatures (Fig. 15). This will be the subject of future study, as it both informs the physical understanding and interpretation of the simulations and observed quasar variability, and presents potential new observational tests and constraints on these accretion disk models, from X-ray through infrared wavelengths. Given the extensive comparisons needed to quantify this, we have focused here on qualitative properties which are robust to short-timescale variations, though note where they are important.

More radical, qualitative “state changes” in the simulations are rare, but we do see them in some cases (illustrated alongside more general variability in the morphology and topology of the magnetic and velocity fields of the inner disk in Fig. 15). Specifically, we have only seen this occur in the inner disk ($\lesssim 200 R_g$), in our highest- \dot{M} (super-critical, $f_{\dot{M}} = 1$,

⁶ Note this is *not* the same as saying the effective or apparent size of the disk or “emitting region” at a given wavelength is $\sim 10 R_g$, because almost all of this luminosity is reprocessed or scattered at larger radii.

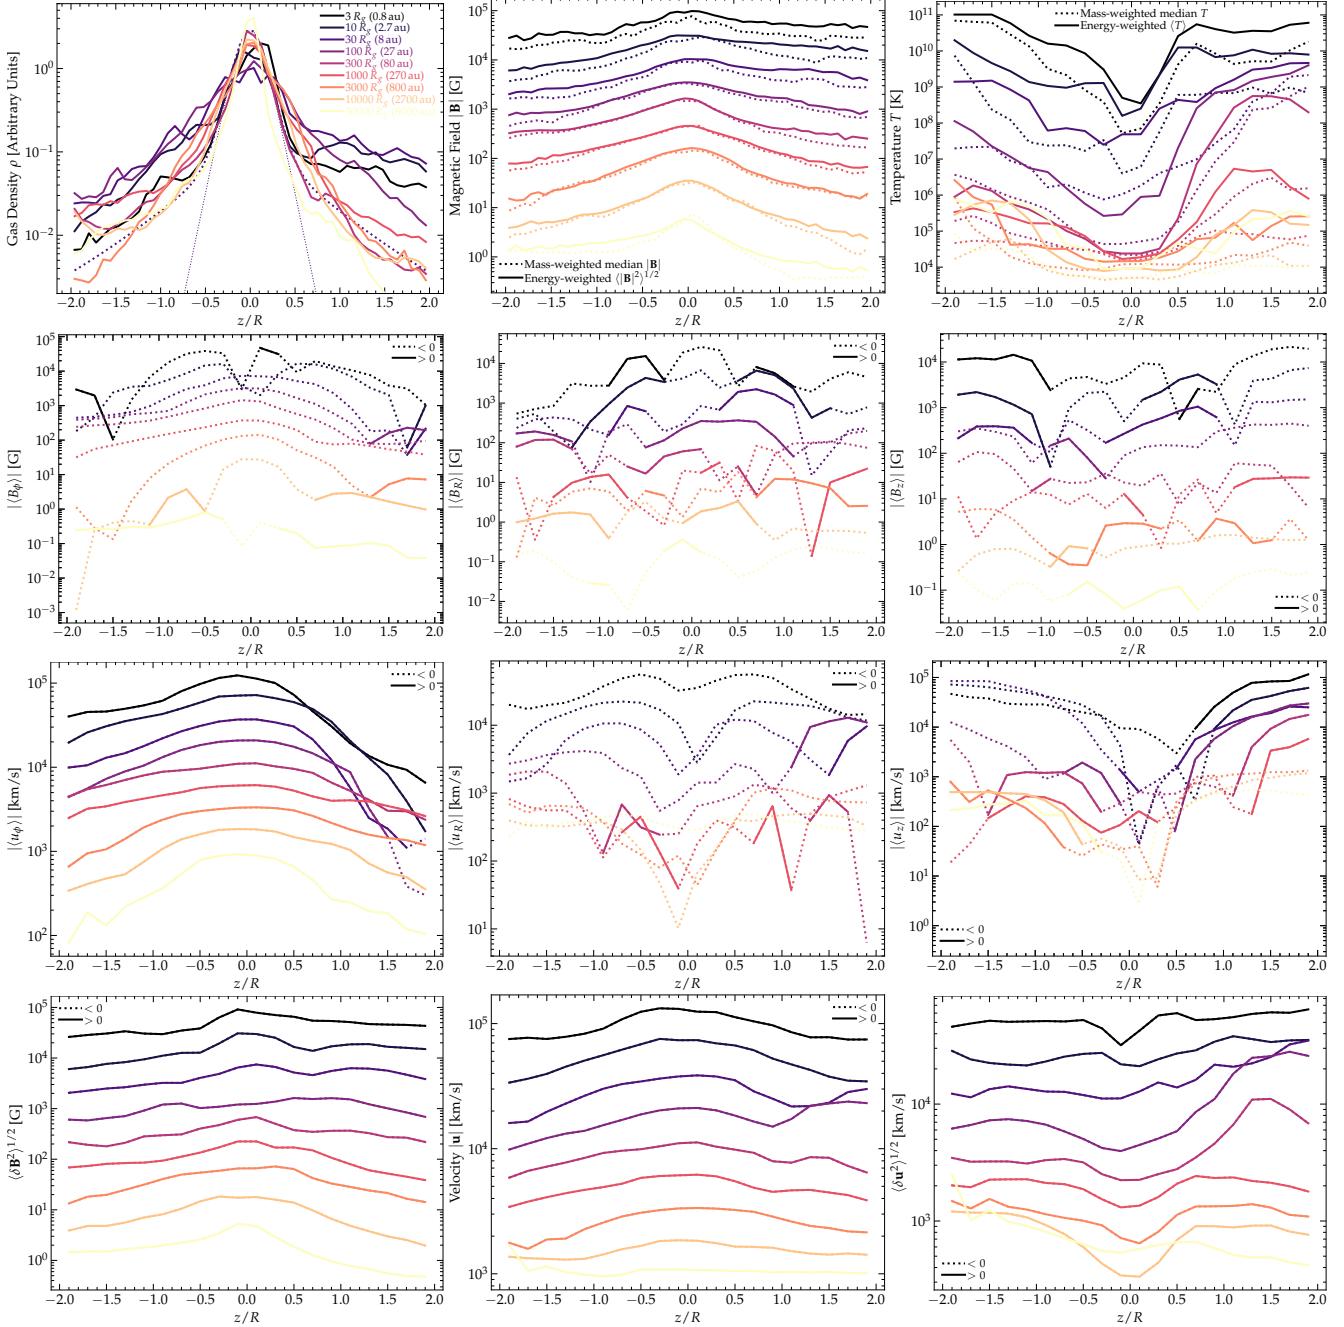


FIG. 14.— Vertical profiles at different cylindrical radii. (1) Density ρ . (2) Magnetic field $|B|$ (primarily B_ϕ). (3) Gas temperature T_{gas} . (4) $\langle B_\phi \rangle$. (5) $\langle B_R \rangle$. (6) $\langle B_z \rangle$. (7) $\langle v_\phi \rangle$. (8) $\langle v_R \rangle$. (9) $\langle v_z \rangle$. (10) $\langle |\delta B|^2 \rangle^{1/2}$. (11) β . (12) $|v|$. (13) $\langle |\delta v|^2 \rangle^{1/2}$.

$\dot{m} \gg 1$) simulation, and only if we refine to the highest level with the inner boundary at $\lesssim 3 R_g$ (as opposed to even at $\sim 15 - 30 R_g$), including our full physics (most importantly radiation pressure forces and MHD; see § 4.10). And seeing such a change typically require running for many dynamical times ($\sim 10^4 - 10^5 \Omega^{-1}$) at the smallest radii. Given all of that, we semi-stochastically see the following series of events. (1) Something (for example, a fluctuation in the accretion rate and disk size/incoming angular momentum direction/warp of the inner disk) makes the outflows somewhat stronger and, more importantly, wider angle and/or able to intercept/entrain some of the inner disk. In our simulations, these outflows begin as radiation-pressure driven (see § 4.7), potentially ex-

plaining the critical role of radiation pressure and zooming into the smallest radii (where the radiation is generated). (2) The outflows then “blow out” a cavity, depleting the central ($\sim 10 R_g$) disk density. (3) The field lines are stretched with the outflow, weakening the mean toroidal and strengthening the coherent poloidal/vertical field, until there is a stronger (volume-averaged) net vertical (seeded by turbulent poloidal loops) flux threading the midplane. (4) The outflow continues propagating, and resembles a “flare” at larger radii. As the outflow propagates, it carries radiation with it from the central disk which had been trapped/advectioned with the accretion flow, as well as blowing-off some surface layers of the disk, lowering the opacity locally and allowing trapped photons to

escape. This produces a propagating “pulse” of photons with higher radiation temperature (potentially observed in EUV or X-rays) and instantaneously high radiative efficiency. Naively, the initial duration of the flare if it escapes would be something like a few $\sim 10^2 - 10^4 G M_{\text{BH}}/c^3$ (hours to months), with a total flare energy typically in the $\sim 10^{51} - 10^{52}$ erg range, but the observed radiation frequencies and durations could be strongly modified by reprocessing further out. (5) The inner accretion disk ($< 100 - 1000 R_g$) is left in a much more MAD-like state (see Kaaz et al. 2024). There is a dominant net vertical field $\langle B_z \rangle$ which has $v_{A,z} \sim 0.2 - 1 v_K$ (can hold up the diffuse gas), with accretion proceeding through this magnetosphere via Rayleigh-Taylor-like interchange instabilities or non-axisymmetric structures (clumps, filaments, etc.). (6) Run long enough, we have a couple examples which eventually have enough mass pile-up and accrete “through” this magnetospheric radius that they rebuild a toroidal-dominated dense disk in the innermost radii, which then fills the magnetosphere from the inside-out, returning the system to a mean toroidal-flux state. But this can require a few orbital times at the radii *outside* of the magnetosphere, which translates to $\gtrsim 10^6 G M_{\text{BH}}/c^3$ (so a very long time relative to near-horizon timescales).

These transitions are broadly similar to those shown in Kaaz et al. (2024), who extend our simulations by taking one of our early snapshots of our highest- \dot{m} ($f_{\dot{M}} = 1$) simulation, and map it to their fully general-relativistic GR-MHD code HAMR (albeit with a more simplified treatment of radiation and thermochemistry). Interestingly, while the transitions above take a long time to occur here, they appear in Kaaz et al. (2024) more rapidly near-horizon, and occur there even in runs without radiation. This may be because the different GR horizon boundaries and other effects produce a strong jet/wind launched from near-horizon scales in Kaaz et al. (2024), even without radiation. This suggests that outflows may play an interesting causal, or at least self-reinforcing, role in such state changes.

Given the restricted conditions/simulations under which we see these changes occur, we focus throughout on the more steady-state (non-burst) behaviors in the simulations. Still, these outbursts are interestingly similar in their timescales, energetics, and radiation properties to observed flares (and potentially some spectral state changes) seen in luminous/high- \dot{m} AGN (see e.g. Boller et al. 1997; Graham et al. 2017, 2020; Trakhtenbrot et al. 2019; Ni et al. 2020; Frederick et al. 2021; Zhu et al. 2021). So they will be studied and compared to observations in detail in future work.

4. DISCUSSION

4.1. Physics of Radiation Pressure Saturation in a Trans/Super-sonically Turbulent Disk

In Paper III and Hopkins (2024), we argued analytically that even at $\dot{m} \gg 1$, the inner disk should feature roughly $P_{\text{rad}} \sim P_{\text{mag}} \sim P_{\text{turb}} \gg P_{\text{thermal}}$, not $P_{\text{rad}} \gg P_{\text{mag}}$. As the inner disk is optically thick, the radiation pressure should be $P_{\text{rad}} \approx (4/3)(\sigma_B/c)T_{\text{mid}}^4$, in terms of the radiation temperature T_{mid} , with effective temperature T_{eff} defined by the surface flux $F_{\text{rad}} = \sigma_B T_{\text{eff}}^4$ and related to T_{mid} in general by some function $T_{\text{mid}}^4 = f_{\tau} T_{\text{eff}}^4$ (f_{τ} defined below). If we assume gravitational energy losses come out in radiation so $\sigma_B T_{\text{eff}}^4 \approx (3/8\pi)\dot{M}\Omega^2$, then $P_{\text{rad}} \approx (4 f_{\tau}/3 c)(3/8\pi)\dot{M}\Omega^2$. Meanwhile these models assume (consistent with the sim-

ulations) trans-Alfvénic turbulence, so magnetic and turbulent pressure scale similarly with $P_{\text{mag+turb}} \equiv P_{\text{mag}} + P_{\text{turb}} \approx \rho(v_A^2 + v_t^2) \approx \Sigma_{\text{gas}} v_t^2 / H$. If Reynolds and/or Maxwell stresses dominate inflow, or equivalently gravitational energy goes into kinetic/magnetic energy before being dissipated radiatively after shocks/reconnection, then $\Sigma_{\text{gas}} v_t^2 \Omega \sim (3/4\pi)\dot{M}\Omega^2$, so $P_{\text{mag+turb}} \sim (2/v_t)(3/8\pi)\dot{M}\Omega^2$. Thus $P_{\text{rad}}/P_{\text{mag+turb}} \sim (4 f_{\tau}/3 c)/(2/v_t) \sim v_t/(c/f_{\tau})$.

So $P_{\text{rad}} \gg P_{\text{mag+turb}}$ requires $c/f_{\tau} \ll v_t$. If radiation transport is dominated by diffusion, then $f_{\tau} \sim \tau$ and $c/f_{\tau} \sim c/\tau$ is the bulk radiation transport speed. But if $c/\tau \ll v_t$, then diffusion cannot dominate radiation transport – vertical transport will be turbulent/advection/conductive, and we must modify f_{τ} . Consider a mixing length-type approximation: the radiation flux in steady state should be vertically constant so $F_{\text{rad}} \sim D_{\text{rad}} \nabla a T_{\text{mid}}^4 \sim \sigma_B T_{\text{eff}}^4$ with D_{rad} the radiation transport coefficient $\approx c/(\kappa \rho)$ for pure diffusion and $\approx v_t \ell_t \approx v_t^2/\Omega$ for turbulent transport with largest eddy turnover time Ω^{-1} . This gives $T_{\text{mid}}^4/T_{\text{eff}}^4 \equiv f_{\tau} \approx c H/D_{\text{rad}}$ which in turn gives the usual $f_{\tau} \approx \tau$ for radiative diffusion, and $f_{\tau} \approx c H \Omega / v_t^2 \approx c/v_t$ for turbulent transport. So when $\tau \lesssim c/v_t$, the transport will be via radiative diffusion with $f_{\tau} \sim \tau$ (provided we still have $\tau \gtrsim 1$), with $P_{\text{rad}} \lesssim P_{\text{mag+turb}}$, and when $\tau \gtrsim c/v_t$, we transition to turbulent vertical transport with $f_{\tau} \sim c/v_t$ and therefore $P_{\text{rad}} \sim P_{\text{mag+turb}}$.

Key to this is that the turbulence is trans or super-sonic. In highly sub-sonic turbulence ($v_t \ll c_s$) in a shearing disk, the maximum turbulent coherence length/driving scale limited by shear is $\sim v_t/\Omega \ll c_s/\Omega \sim H$. This would mean that the turbulence would be confined to a midplane layer with height $\ll H$, so it is not even clear if we could safely assume eddies actually could carry radiation all the way to the photosphere at $|z| \sim H$. But even if we assume such turbulence is volume-filling in the disk with said coherence length, in the derivation above in the turbulence-dominated limit $f_{\tau} \rightarrow c H \Omega / v_t^2 \rightarrow (c_s/v_t)(c/v_t)$, and so (1) the turbulent-transport-dominated limit is much harder to achieve, in practice, and (2) even in the said limit, we would have $P_{\text{rad}}/P_{\text{mag+turb}} \sim v_t/(c/f_{\tau}) \sim c_s/v_t \gg 1$. Thus, for *subsonic* turbulence, disks can become radiation-pressure dominated in the manner envisioned in e.g. SS73. But for a *trans/supersonically* turbulent disk, as inevitably occurs when $\beta_{\text{thermal}} \ll 1$ in the disk, then the disk self-regulates to $P_{\text{rad}} \sim P_{\text{mag+turb}}$, and never truly becomes radiation-pressure dominated. Though the radiation pressure is, by definition, non-negligible here.

Note that saturation with $P_{\text{rad}} \sim P_{\text{mag,turb}}$ has also been widely seen in more idealized simulations of supercritical accretion flows beginning from a thin, thermal-pressure dominated SS73-like initial condition (Jiang et al. 2019; Davis & Tchekhovskoy 2020; Utsumi et al. 2022; Huang et al. 2023). There, the “direction” in which this solution is approached is somewhat opposite to that here, as the radiation pressure puffs up the disk and drives stronger turbulence which leads to a turbulent dynamo saturating the turbulent $P_{\text{mag,turb}}$ near P_{rad} . But it suggests this is a robust state. The major difference is that these systems also have a large midplane mean-field component, so can saturate with the mean/total magnetic pressure somewhat larger than P_{rad} .

4.2. Physics of the Radiative Efficiencies

For similar reasons, we expect the inner disk to be radiatively efficient for $0.01 \lesssim \dot{m} \lesssim 1$, and to become increasingly radiatively inefficient at $\dot{m} \gg 1$. As discussed in Hopkins (2024), even with the turbulent mixing of radiation in § 4.1, because the opacities do not discontinuously transition to $\tau \ll 1$ at $|z| \sim H$ (the coherence length of the turbulence) – i.e. there is no sharp “edge” of the disk and H is not dramatically smaller than \dot{R} – the turbulence should not qualitatively change the ability of radiation to escape in the radiatively saturated regime of § 4.1. There, the diffusive radiation escape speed is $\sim c/\tau_{\text{eff}}$ (where τ_{eff} can reflect the minimum τ seen over the eddy lifetime in the turbulence-dominated regime), which should become smaller in our highest-accretion-rate run ($f_M = 1$, $\dot{m} \gtrsim 20$) than the advective speed v_R at $R \lesssim 300 - 1000 R_g$ according to Figs. 5 & 8. This corresponds well to where we measure the advective flux becoming larger than the comoving radiation flux in Fig. 11. This is also consistent with the radii predicted for this transition from the analytic disk models in Hopkins (2024).

The major difference between this and classic advection-dominated “slim disk” scenario (Paczyński & Wiita 1980; Abramowicz et al. 1988) is that those models and other like them assume (as did SS73) that $\beta \gg 1$ throughout the midplane, which means that they only account for advection of thermal energy. Here, we plainly see from Fig. 5, and can directly check in Fig. 11, that most of the advected energy is magnetic, particularly in the form of the dominant mean toroidal field. This modifies the solutions as discussed below.

On the other hand, for our $\dot{m} \ll 1$ simulations we see in Fig. 16 that the disk is radiatively efficient, with radiation escape times $\sim r/(c/\tau_{\text{eff}})$ which are fast because of the modest $\tau \sim 1$ in the disk; specifically, faster than accretion timescales down to the ISCO ($\sim r/|\langle v_r \rangle| \sim 10 \Omega^{-1}$). This is expected from the arguments above and models in (Hopkins 2024).

4.3. Dependence on Accretion Rate

We see in Figs. 16 & 17 that the disk *structural* properties scale with \dot{m} , to leading order, quite close to the similarity solution predictions in Paper III. Recall the simulations have fixed M_{BH} and are compared at the same radii and times, but different (rescaled) initial gas mass. The initial conditions are rescaled versions of each other, but all begin with no gas interior to $< 300 R_g$, so predictions interior are entirely distinct and at larger radii we see as we evolve the runs they reach different local equilibria (only at the very largest R do they reflect the initial conditions, because the simulation time is a small fraction of Ω^{-1} at e.g. $\gg 10^5 R_g$). To leading order, we see the gas densities and kinetic, magnetic, and radiation energy densities/pressures, and absolute magnitude of the stresses, scale $\propto \dot{m}$ as predicted. This means dimensionless quantities like β , the accretion timescales $t_{\text{acc}}/t_{\text{orbit}}$, Alfvén Mach numbers M_A , and ratio of pressures/stresses in different components, and density/magnetic/turbulent scale heights H/R are approximately independent of \dot{m} . This is already highly non-trivial: it implies that lower accretion rates are, to first order, simply a result of lower mass supply, not mass moving “less quickly” through the disk. Moreover it contrasts strongly with the predictions of standard α disk models like SS73, which predicts $\rho \propto \dot{m}^{-2}$ (or $\propto \dot{m}^{0.4}$) in the inner (outer) regions of these disks (their “region (a)” and “region (b) or (c)”, respectively); or $H/R \propto \dot{m}^{0.2-1}$, $\delta v^2/v_K^2 \propto \dot{m}^{0-0.85}$, $t_{\text{acc}}/t_{\text{orbit}} \propto \dot{m}^{-2}$ to $\dot{m}^{0.4}$, etc.

Per above (§ 4.2), the luminosity and radiative efficiency

also scale with \dot{m} , roughly as expected though with the important aspect that advection makes the radiative efficiency not self-similar (as seen in other simulations, see e.g. Jiang & Dai 2024 for a review). Correspondingly weaker radiation-pressure driven outflows and $a_{\text{rad}}/a_{\text{grav}}$ at lower \dot{m} (discussed below, § 4.7) are expected.

The biggest systematic deviations from self-similarity are in the thermal properties (temperature, species abundances, opacities), which should not be scale-free. Still these roughly follow the predictions from more detailed analytic calculations in Hopkins (2024) of their properties in a disk whose structural properties follow the similarity solutions of Paper III, and their variation is relatively weak with \dot{m} (recall, \dot{m} varies by a factor of ~ 1000 between the different runs). The dust sublimation radius moves out at higher luminosity or \dot{m} , but atomic and molecular gas can actually persist in the midplane further interior (at smaller R) at higher \dot{m} owing to the higher densities of the midplane leading to both much stronger self-shielding, but also much higher recombination/molecule formation rates. This in turn non-linearly modifies the opacities (though the total optical depth, $\propto \kappa \Sigma_{\text{gas}}$, varies more with Σ_{gas} which varies roughly as $\propto \dot{m}$ here, while the opacities κ vary relatively weakly). As predicted in Hopkins (2024, see their Figs. 6, 12, & 13), the lower- \dot{m} disks are hotter in terms of their volume-filling T_{gas} (note this is opposite the thermal-pressure dominated disk prediction, for which $T \propto \dot{m}^0$ to $\dot{m}^{0.4}$), and have more very hot (quasi-virialized) atmosphere/coronal gas, out to larger radii, even though the dust sublimation radius moves inwards. The behavior of the radiation temperature in Fig. 16 needs a bit more interpretation, however: we discuss this in more detail in § 4.9, where we show it owes in part to highly non-LTE effects, where the role of optically-thin cooling and radiation becomes more important at lower \dot{m} .

4.4. Analytic Models & Modified Radial Scalings of Structural Properties

Hopkins (2024) argued that at $\dot{m} \gg 1$, the flow should become advection-dominated with most of the advected energy in the form of the mean toroidal field (as we see in Figs. 5 & 11). They noted that, akin to Narayan & Yi (1995); Narayan et al. (1998), there is a simple similarity solution for this regime which was already discussed in Paper III, namely that interior to the radius where advection dominates, the flow would be modified to:

$$\frac{v_{\text{turb}}}{v_c} \sim \frac{v_A}{v_c} \sim \frac{H}{R} \rightarrow \psi \sim 0.2, \quad (1)$$

$$\Sigma_{\text{gas}} \rightarrow \frac{\dot{M} \tilde{g}(x_g)}{2\pi\psi^2 \sqrt{GM_{\text{BH}}R}} \quad (2)$$

$$\sim 4 \times 10^7 \frac{M_{\odot}}{\text{pc}^{-2}} \left(\frac{\dot{M}}{3M_{\odot} \text{ yr}^{-1}} \right) \sqrt{\frac{1 - x_g^{-1}}{x_g/100}} \left(1 - \sqrt{\frac{3}{x_g}} \right).$$

The similarity parameter ψ is set by continuity between the radii $R \lesssim R_{\text{adv}}$ where the solution becomes advective, and the solutions at larger scales to be $\psi \approx (R_{\text{adv}}/r_{\text{ff}})^{1/6} \sim 0.2 (m/100)^{1/6}$ for the parameters here. we have incorporated the correction factor $\tilde{g}(x) \equiv \sqrt{1 - 1/x} (1 - \sqrt{x_{\text{in}}/x})$ which accounts for both the shape of the PW potential (the first term) and assumption that the angular momentum vanishes at some x_{in} as in SS73 (who took $x_{\text{in}} = 3$ for the ISCO of a non-rotating BH).

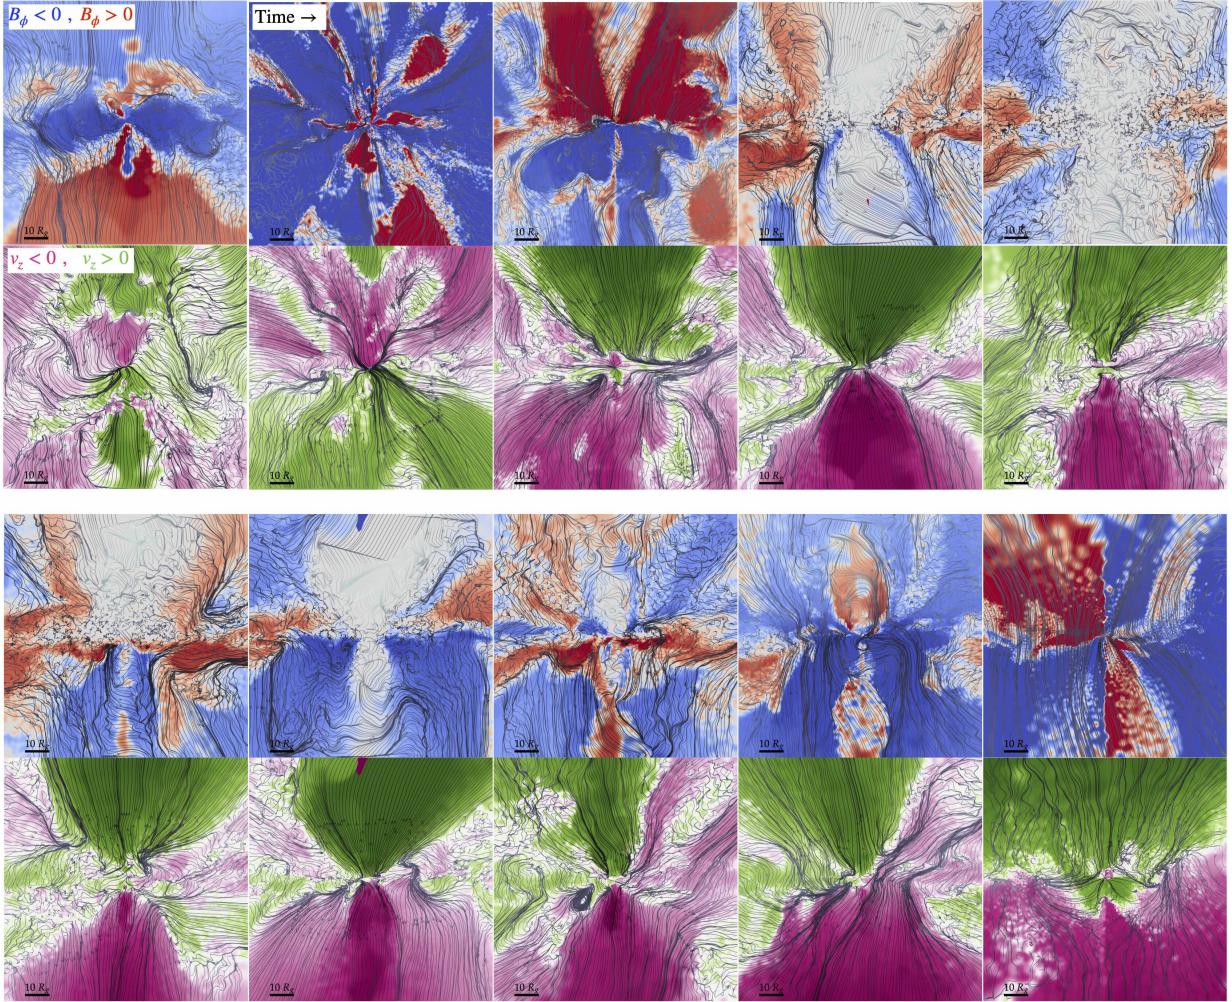


FIG. 15.— Time series of edge-on maps of the magnetic field (showing sign of B_ϕ , as Fig. 7) and velocity (sign of v_z , as Fig. 9), in the inner disk ($< 100 R_g$) of our highest- \dot{m} (supercritical; $f_{\dot{M}} = 1$) simulation. Snapshots span a timescale $\sim 10^6 GM_{\text{BH}}/c^3 \sim 2$ years, but not uniformly (chosen to highlight different behaviors rather than quantitative time evolution). As discussed in § 3.5, the detailed morphology can vary considerably, and there are even times of inflow and times where the symmetry properties of the toroidal field change from even about the midplane to odd, sometimes associated with flares and strong bursts of outflow which can disrupt the inner disk. These are discussed in Kaaz et al. (2024), and the time variability of these simulations will be studied in future work.

If we account for the fact that \dot{M} is not actually exactly constant with radius at the time studied here (Fig. 5), we see in Figs. 5–11 that these equations, with constant $\psi \sim 0.2$, provide a remarkably accurate prediction of the disk dynamical and mass and density structure interior to the radius where the system becomes advective. Furthermore we see in Fig. 16 that these analytic solutions explain the zeroth-order qualitative scalings with both \dot{m} and r in the simulations.

The similarity models do not perfectly reproduce the chaotic, non-linear complexity of the simulations, and the deviations can be significant (e.g. up to ~ 1 dex in midplane n , Fig. 16). As discussed above, these deviations are often “quasi-steady-state” in the sense that they owe to global structures/modes (e.g. spiral arms, gaps, rings, interactions between outflows and inflows) that break the self-similarity and can persist for many orbital times at a given radius (as opposed to local turbulent fluctuations which have coherence times $\lesssim \Omega^{-1}$). So they cannot be trivially “averaged out” (especially since any time average here at large R would average over many different local steady-state configurations at small R). But it is worth contrasting these magnetically-dominated disk models with a standard (assumed $\beta_{\text{thermal}} \gg 1$) α -disk

model as an alternative similarity solution, as presented in detail in Paper III and Hopkins (2024). Per § 4.3, the α -disk model predicts a qualitatively different (sometimes opposite) dependence on \dot{m} of key quantities like n , H/R , δv_{turb} , t_{acc} , T_{gas} ; likewise as shown in Hopkins (2024) the qualitative dependence on r can be different for quantities like H/R and τ . But the most notable difference is in the characteristic values: α -disk models (with the same \dot{M} , M_{BH} , and R simulated here) predict H/R systematically smaller by $\sim 2 - 3$ dex, τ and Σ_{gas} larger by $\sim 4 - 6$ dex, $|\mathbf{B}|$ larger by $\sim 1 - 2$ dex, and midplane densities n larger by $\sim 6 - 9$ dex (Hopkins 2024).

4.5. Extended Scattering Structures & Effective Size of the Emission Regions

The extended scattering structures visible in Figs. 12–13 are a natural consequence of the large scale-height of the disk (Figs. 2, 5, 13, 14), and correspondingly relatively slow falloff of $\rho(z)$ in Fig. 14. This is discussed in more detail for the analytic models in Hopkins (2024), but basically since $H/R \sim 0.1 - 1$ at the radii of interest here, (1) a large fraction of the radiation from the inner disk directly intercepts the outer disk, and (2) the density at $|z| \sim R$ cannot be vastly smaller than

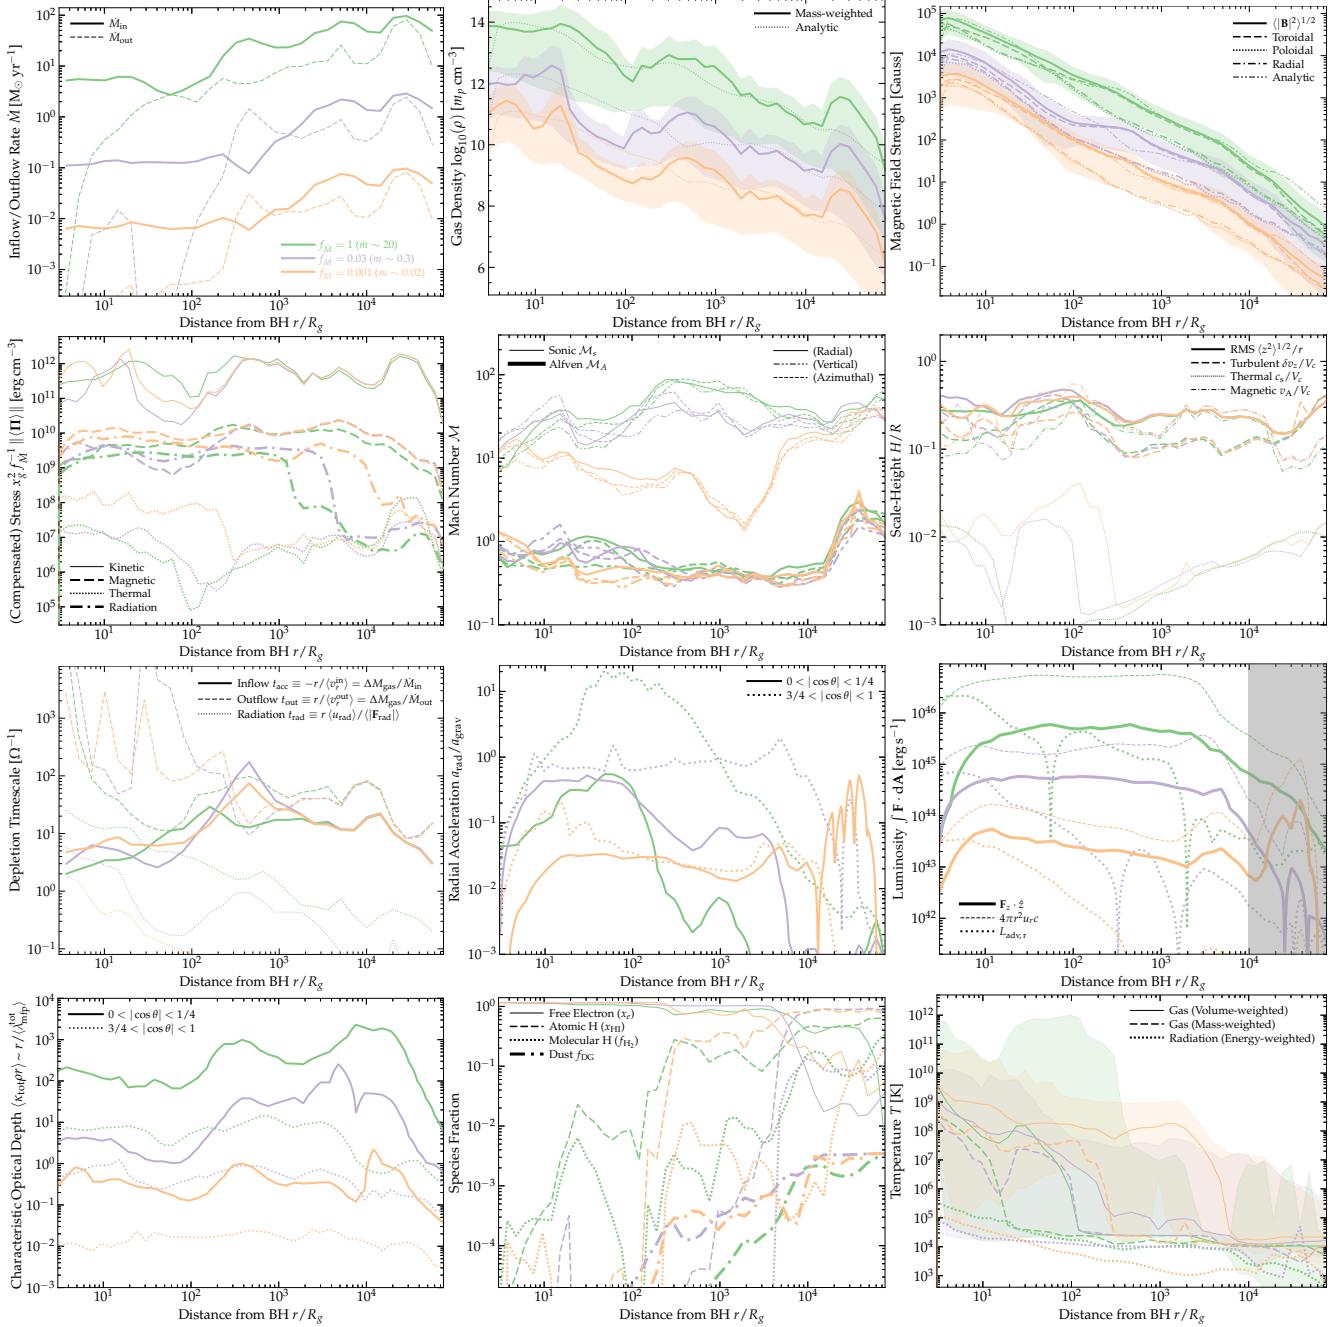


Fig. 16.— Radial profiles as Figs. 5, 6, 8, 11, but comparing our three simulations with different initial gas accretion rates (all else fixed): $f_M = (1, 0.03, 0.001)$ or $\dot{m} \sim (20, 0.3, 0.02)$ (green, purple, orange, respectively). To leading order, properties like \dot{M} , ρ , $|\mathbf{B}|$, Σ_{gas} or τ , and different energy densities scale with \dot{m} while dimensionless quantities like H/R , \mathcal{M}_A , $t_{\text{acc}}/t_{\text{orbit}}$ are independent of \dot{m} , as predicted by simple magnetically-dominated disk similarity solutions (Paper III). While thermochemistry is not scale-free, $\beta_{\text{therm}} \ll 1$ at all radii, with $\beta_{\text{tot}} \sim 0.1 - 1$ in the inner disk, in all cases. Below $\dot{m} < 0.02$, there ceases to be an optically-thick accretion disk. Continuum radiation pressure ($\alpha_{\text{rad}}/\alpha_{\text{grav}}$) becomes weaker at lower \dot{m} , and radiative efficiencies decrease with \dot{m} as $\epsilon_r \sim (0.15, 0.1, 0.02)$ for $\dot{m} \sim (0.02, 0.3, 20)$.

the midplane density (since even for an exponential vertical profile, this is just a few e -foldings $|z|/H \sim 1 - 10$). The latter means that an extended scattering halo/atmosphere in more polar directions is automatically expected, even without any “additional” component like outflows or jets. That directly leads to the curved flux lines seen in Fig. 12.

Correspondingly, from Fig. 11, we see that the radiation from the inner disk has an $\mathcal{O}(1)$ scattering component (but primarily scattering at these larger radii rather than thermal reprocessing), out to radii of a few thousand R_g or $\sim 10^{16}$ cm.

This is intriguingly similar to the observed sizes of the optical/NUV emission observed in microlensing (Dai et al. 2010; Blackburne et al. 2011; Jiménez-Vicente et al. 2014; Cornachione & Morgan 2020) and reverberation mapping (Cai & Wang 2023; Ren et al. 2024) studies. The fact that these studies also see sizes independently of wavelength has led to widespread speculation that this must owe a process like Thompson scattering (being achromatic) from an extended surface – but the historical challenge has been that an $\mathcal{O}(1)$ fraction of the light must be reprocessed, which is not pos-

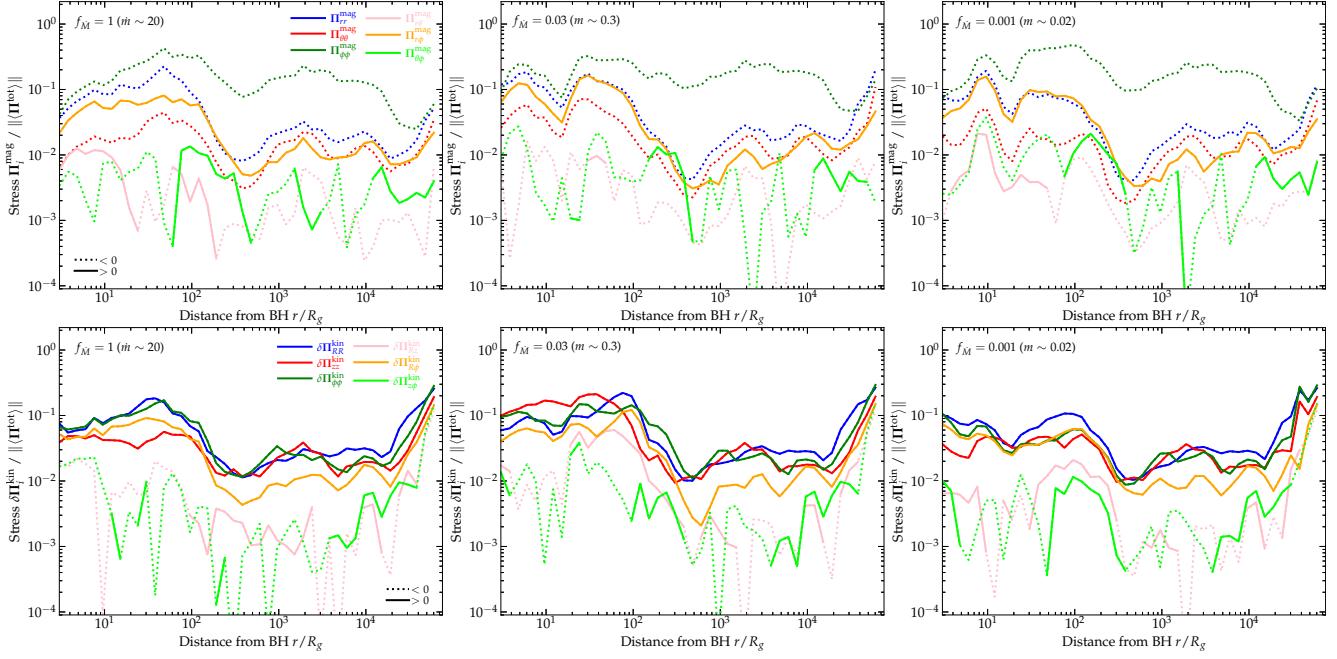


FIG. 17.— Radial profiles as of Reynolds (top; as Fig. 8) & Maxwell (bottom; as Fig. 6) stresses, comparing our different accretion-rate simulations $f_{\dot{M}} = (1, 0.03, 0.001)$ or $\dot{m} \sim (20, 0.3, 0.02)$ (left, middle, right). The dimensionless scale of the stress is self-similar, and in all cases angular momentum transport is dominated by the $R\phi$ stress (comparable to the total stress) with mean Maxwell marginally larger than fluctuating Maxwell and/or Reynolds in the inner disk.

sible in most physical models of SS73-like α disks (with $H/R \sim 0.001 - 0.01$ at these radii) and/or winds from the disk surface (Czerny et al. 2003; Dai et al. 2010; Kamraj et al. 2022; Tortosa et al. 2022). In future work, it therefore seems important to forward-model these observations directly.

4.6. Effective Temperature of the Emerging Radiation Field

In a disk which is radiatively efficient at all radii and very thin ($H \ll R$), so radiation from the surface escapes to infinity, we expect the effective radiation temperature of the radiation *emitted* at each radius/annulus R to scale as $T_{\text{rad, eff}}^{\text{em}} \sim (3 \dot{M} \Omega^2 / 8\pi\sigma_B)^{1/4} \sim 15000 \text{ K} (\dot{M} / 10 \text{ M}_\odot \text{ yr}^{-1})^{1/4} (R / 1000 R_g)^{-3/4}$ (scaling to the BH mass here). At the opposite extreme however, if the system is radiatively inefficient with most of the radiation coming from small R , and the radiation is reprocessed quasi-spherically (i.e. $H \sim R$ and the system remains optically thick at large heights), then $T_{\text{rad, eff}}^{\text{em}} \sim (L / 4\pi r^2 \sigma_B)^{1/4} \sim 15000 \text{ K} (L / 6 \times 10^{44} \text{ erg s}^{-1})^{1/4} (R / 1000 R_g)^{-1/2}$. For purposes of observational probes like the polarization studies in Kishimoto et al. (2008), even these opposite extremes give very similar predictions for the temperature initially emitted at different radii, and both are consistent with the “un-processed” spectral shapes seen therein.

However, what can differ more dramatically between models is not the temperature initially emitted at each R , but the effective temperature of the integrated spectrum (including reprocessing) escaping to infinity, $T_{\text{rad, eff}}^{\infty}$. In the thin, radiatively-efficient disk model with no reprocessing, the integrated spectrum is dominated by light from the innermost radii (around the ISCO), with $T_{\text{rad, eff}}^{\infty} \sim 10^6 \text{ K}$ predicted for the BH mass and accretion rate here. On the other hand, quasi-spherical reprocessing (from a disk/photosphere with appreciable H/R) acts like a photosphere, so the escaping spectrum has $T_{\text{rad, eff}}^{\infty}$

determined to first order by the radiation temperature at the radii where the absorption optical depth drops below unity – roughly $R \sim 500 R_g$ in our simulations here (Fig. 11). The latter predicts $T_{\text{rad, eff}}^{\infty} \sim 10,000 - 30,000 \text{ K}$ for our simulations (in good agreement with the analytic predictions in Hopkins 2024, who quote $T_{\text{rad, eff}}^{\infty} \sim 4 \times 10^4 \text{ K} \epsilon_{r, 0.1}^{11/32} m_7^{-3/64} \dot{m}^{-7/32} \tilde{Z}^{-3/16} \sim 2 \times 10^4 \text{ K}$ for parameters simulated here).⁷ This is much closer (compared to the thin-disk $T_{\text{rad, eff}}^{\infty} \sim 10^6 \text{ K}$ prediction) to what we see actually emerging in the optically-thin polar cone in the simulations, once the radiation is escaping (Figs. 4, 5, 12, 13, 16). It is also much closer to the actual effective temperatures observed (in total light) in quasar SEDs (Peterson 1997; Vanden Berk et al. 2001; Richards et al. 2006; Shen et al. 2020; Cai & Wang 2023). More detailed spectral predictions will be the subject of future work utilizing extensive radiative-transfer post-processing.

4.7. On the Properties of Outflows (Absent Jets and Line Driving)

In all our runs, we can plainly see there there is a non-negligible amount of gas with large radial and/or vertical velocities sitting at large displacements from the midplane, i.e. with $v_r \gtrsim 0.1 - 1 v_K$ and $|z| \gtrsim 0.1 - 1 R$, at a wide range of radii R from $\sim 10 - 10^5 R_g$ (see Figs. 8, 9, 14, 18, 19). These may even appear the same as large-scale winds in the observable spectrum (this will be the subject of future work). However, given the highly supersonic turbulent motions with $\langle \delta v_z^2 \rangle^{1/2} \sim 0.1 v_K$ and corresponding thick disk

⁷ Again note as discussed in § 3.3 that the apparent sudden drop of T_{rad} at the dust sublimation radius to temperatures $\ll 10^4 \text{ K}$ owes to the fact that in-code we evolve a photon-number-weighted effective temperature, so as soon as the dust absorbs and re-radiates even a fraction of the bolometric luminosity, this occurs. But really this simply indicates where the spectrum becomes multi-component with escaping UV/optical and dust emission in the IR.

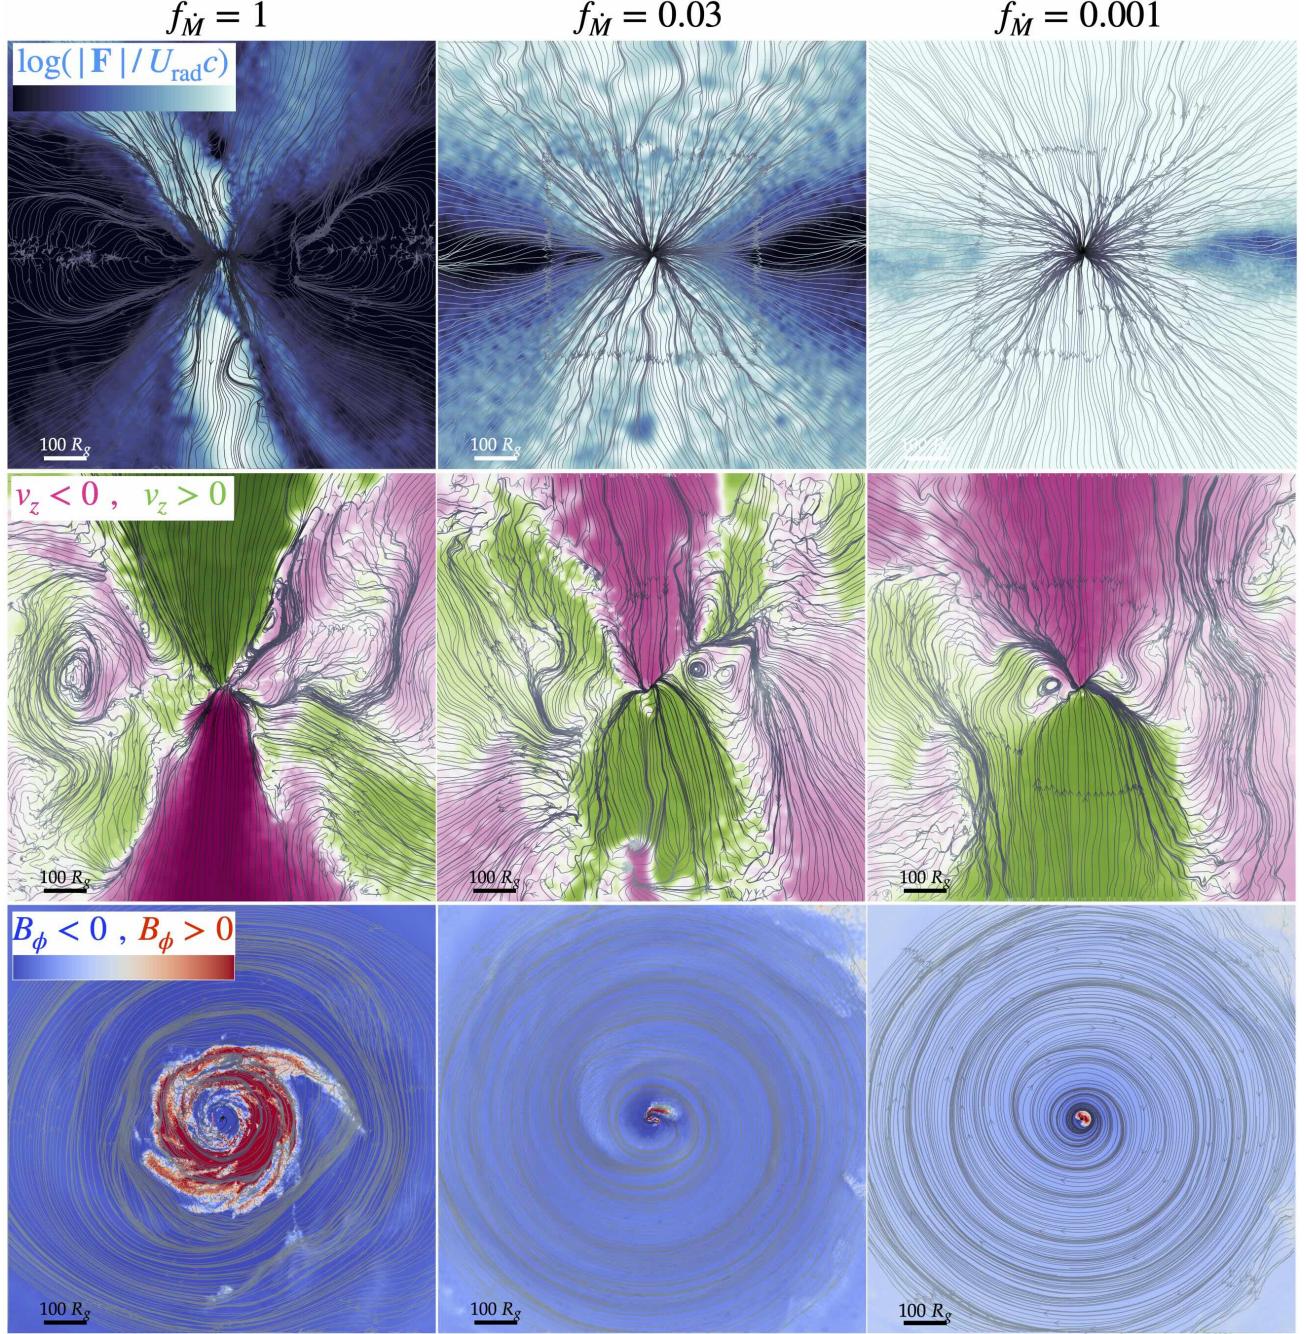


FIG. 18.— Field-line maps comparing runs with different accretion rates $f_{\dot{M}} = (1, 0.03, 0.001)$ or $\dot{m} \sim (20, 0.3, 0.02)$ (left, middle, right) at fixed spatial scale. *Top:* Radiation flux edge-on as Fig. 12. Radiation escapes in a widening bicone at lower \dot{m} , with less scattering so the flux approaches the free-streaming limit ($|\mathbf{F}| \rightarrow u_{\text{rad}} c$). The reduced scattering at low \dot{m} means we see more-radial flux streamlines without as much curvature “into” the disk (from scattering and absorption/reprocessing of inner disk photons by outer disk gas). *Middle:* Velocity field edge-on as Fig. 9. All exhibit weak large-scale coherent outflows, but intermediate-scale semi-coherent outflows rising to $|z|/R \sim 0.1 - 1 \sim H$, i.e. potential “failed winds” or “fountain flows” become less prominent at lower \dot{m} . *Bottom:* Magnetic field face-on as Fig. 7. The toroidal field is dominant in all cases with coherence lengths $\gtrsim R$.

with $H/R \sim 0.1 - 0.3$ (Fig. 5), it is sometimes difficult to tell if these are “failed winds” or “fountain flows” or simply part of the turbulence.

What we do not always see is coherent, stable, steady-state outflows filling most of the volume vertically “above” and “below” the disk (“rising” directly from the disk vertically even in the outer disk at $\gg 100 R_g$), as one might expect in simple analytic models for magnetocentrifugal quasar winds. But we do see some narrower opening-angle “broad jet-like” outflows

driven by radiation pressure in our highest- \dot{m} simulations.

Note that we do *not* resolve the true jet-launching region nor include the physics necessary to properly model relativistic-jet formation,⁸ so work with GRMHD simulations extending these to true horizon scales is needed (presented in Kaaz et al.

⁸ Our Newtonian+PW method does not capture spin or other GR effects, but also treats the inner boundary as a sink domain (into which quantities like magnetic flux can simply advect and be removed from the active grid) in standard fashion, as opposed to a true horizon (which can be important the buildup of near-horizon poloidal flux; see Kaaz et al. 2024).

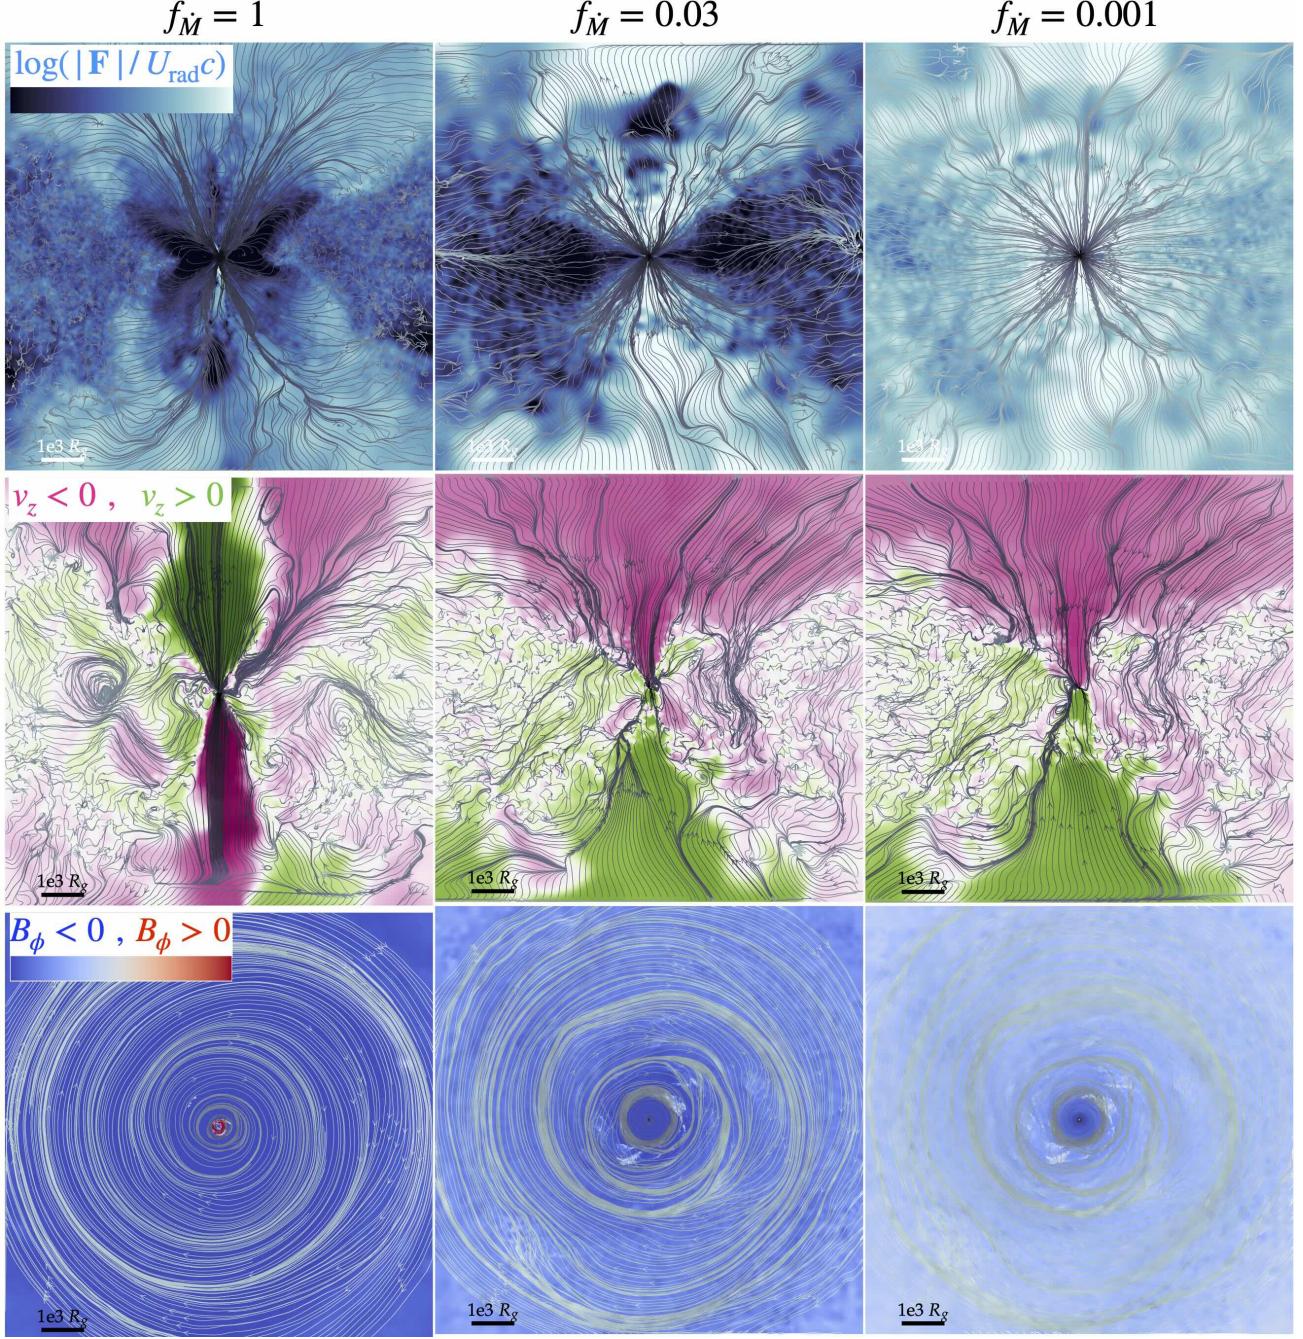


FIG. 19.— As Fig. 18, at 10 times larger scale.

2024). We also do not model radiation line-driving, so cannot properly represent line driven outflows from the broad-line region (one of the most popular candidates for quasar wind acceleration; see e.g. Proga et al. 2000). But simpler, continuum radiation pressure or magnetocentrifugal pressure-driven outflows (at least in the inner, dust-free disk which we evolve for many dynamical times) are captured in our simulations.

4.7.1. Magnetocentrifugal Outflows: Failed Winds, But Potential Drivers of Thick Disks and Trans-Alfvénic Turbulence

In the outer disk, as in Paper II, we see only weak outflow along occasional flow lines coming directly off the disk surface at large radii. Most of the apparent outflow in Figs. 5, 16 owes to radial non-circular motions, and the polar out-

flow at largest \dot{m} is launched from relatively small radii. This implies magnetocentrifugal or buoyancy-driven outflows are weak through most of the disk. That may at first seem surprising in both cases given the magnetically-dominated disk, but should not be. Vertical buoyancy (e.g. Parker-like) modes can drive convective vertical turbulence (and may indeed play an important role in the turbulence here and even in sustaining the toroidal fields; see Johansen & Levin 2008; Gaburov et al. 2012; Squire et al. 2024), but induce relatively little outflow/overshoot above the density scale-height. Magnetocentrifugal outflows require a strong mean vertical field threading the disk to wind up into a (toroidal-vertical) magnetic tower (Spruit 1996; Bai et al. 2016). Indeed wind escape

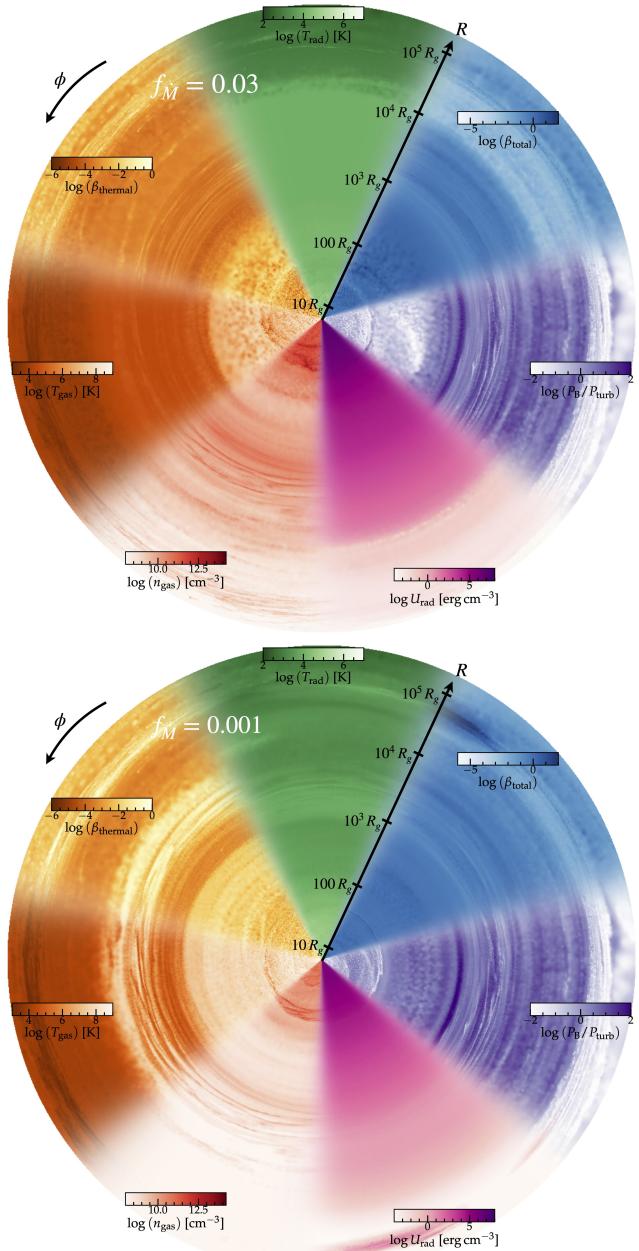


FIG. 20.— Face-on projections showing the entire (log-scaled) dynamic range of the simulations with different accretion rates $f_M = (0.03, 0.001)$ or $\dot{m} \sim (0.3, 0.02)$ (top, bottom), as Fig. 4. Mass and radiation density decrease with lower \dot{m} as expected, while β increases weakly, and \mathcal{M}_A remains similar.

in classic [Blandford & Payne \(1982\)](#)-type models assumes a vertical coherence length of field lines $\ell_{B_z} \rightarrow \infty$, while here we see the vertical coherence scale is $\sim H$. And by the criterion of [Spruit \(1996\)](#), the wind “mass loading” parameter $\mu_{\text{wind}} = \Omega_0 R_0 v_{p,0} 4\pi \rho_0 / B_p^2 \sim v_K / |\langle v_{A,z} \rangle| \sim 100 - 1000$ from Fig. 6 is extremely large, which implies a “failed wind” is likely as the wind terminal velocity is predicted to be $\ll v_A \ll (H/R) v_K$.

So almost any turbulence or exterior ram pressure from gas above/below the disk, or other instabilities would stall or disrupt the escape of said winds to $|z| \gg H$. Indeed as discussed in [Spruit \(1996\)](#); [Ouyed & Pudritz \(1999\)](#); [Anderson et al. \(2005\)](#), when $\mu_{\text{wind}} \gg 1$, especially if $|\langle B_\phi \rangle| \gg |\langle B_z \rangle|$

and/or $c_s \ll v_{A,\phi}$ (all of which are true here), the “wind” stalls or becomes confined within or just above the disk (at $|z| \sim H$). In-plane, in those studies, it then behaves more like one of many sources of weak oscillatory in-plane spiral motions. Out of plane, in the same studies, it continues to drive material up to $|z| \sim H$ then recycles in a failed wind or fountain flow. So although this will not lead to a true wind, it suggests that “failed” magnetocentrifugal or Blandford-Payne type flows could be an important source of the observed trans-Alfvénic vertical turbulence which is critical to maintain the disk scale height at large $H/R \sim 0.1 - 1$.

4.7.2. (Continuum) Radiation Pressure-Driven Outflows

Regarding radiation pressure, recall that we include continuum opacities but not line-driving, so cannot capture line-driven winds which may be extremely relevant given the large opening-angle, irradiated disk surface BLR-like regions (Fig. 12). Note also that we only evolve the simulations long enough for the radiation front to reach equilibrium out to $\sim (0.5 - 5) \times 10^4 R_g$ (depending on whether it is free-streaming or slowed by scattering), once we reach our full refinement around the ISCO. But since most of the radiation is emitted at these small radii, this means that the full effects of radiation-pressure driven outflows at *larger* radii, most notably outside the dust sublimation radius (as large as $\sim 10^5 - 10^6 R_g$ for the most optically-thin polar sightlines) cannot be modeled because we cannot run these simulations long enough. Instead such outflows are better studied in simulations which do not explicitly simulate the gas flows near the horizon (as we do here, imposing a steep timestep penalty), but inject radiation with a sub-grid model for the un-resolved accretion disk at some much larger inner boundary ($\gtrsim 1000 R_g$). So we stress that the absence of strong radiation-pressure driven outflows from either the BLR or torus is an artifact of our numerical limitations here, and not a prediction. In future work, we will model these regions with detailed radiative transfer post-processing in order to better understand wind-launching conditions therein (Bardati et al., in prep.).

With that in mind, consider continuum radiation pressure *within* the accretion disk at radii $\lesssim 10^4 R_g$, where our simulations are more predictive. It is not surprising that our lowest $\dot{m} \sim 0.02$ ($f_{\dot{M}} = 0.001$) simulation does not exhibit strong radiation-pressure driven outflows. It is a factor of ~ 50 below the electron-scattering Eddington limit in luminosity and we confirm the local radiative acceleration a_{rad} (Fig. 16) is similarly below the strength of gravity a_{grav} , as expected. Likewise the intermediate $\dot{m} \sim 0.3$ ($f_{\dot{M}} = 0.03$) case is more marginal, but still has $a_{\text{rad}}/a_{\text{grav}} \sim 0.3$ as expected through the volume-filling region, with $a_{\text{rad}} \ll a_{\text{grav}}$ in most of the disk except in the very innermost disk ($\lesssim 100 R_g$) where this ratio still remains sub-unity. In this latter case there could easily be appreciable radiation pressure driven outflows beginning to appear at $\gtrsim 10^4 R_g$ from radiation pressure on dust, but again the evolution time of the simulation is only just long enough for the body of the radiation (emitted near the ISCO, so requiring our full refinement) to reach these radii. Just like with magnetocentrifugal outflows (§ 4.7.1), radiation pressure driving in Fig. 11 appears sufficient to accelerate material into non-trivial “fountain” or “failed wind” motions with $|\delta v| \sim 0.1 - 0.3 v_K$, comparable to the turbulent speeds, so some radiation-pressure-driven flows are present and could contribute to the observed velocity structure of lines and disk thickness in this ($\dot{m} = 0.3$) simulation, but they are not creating

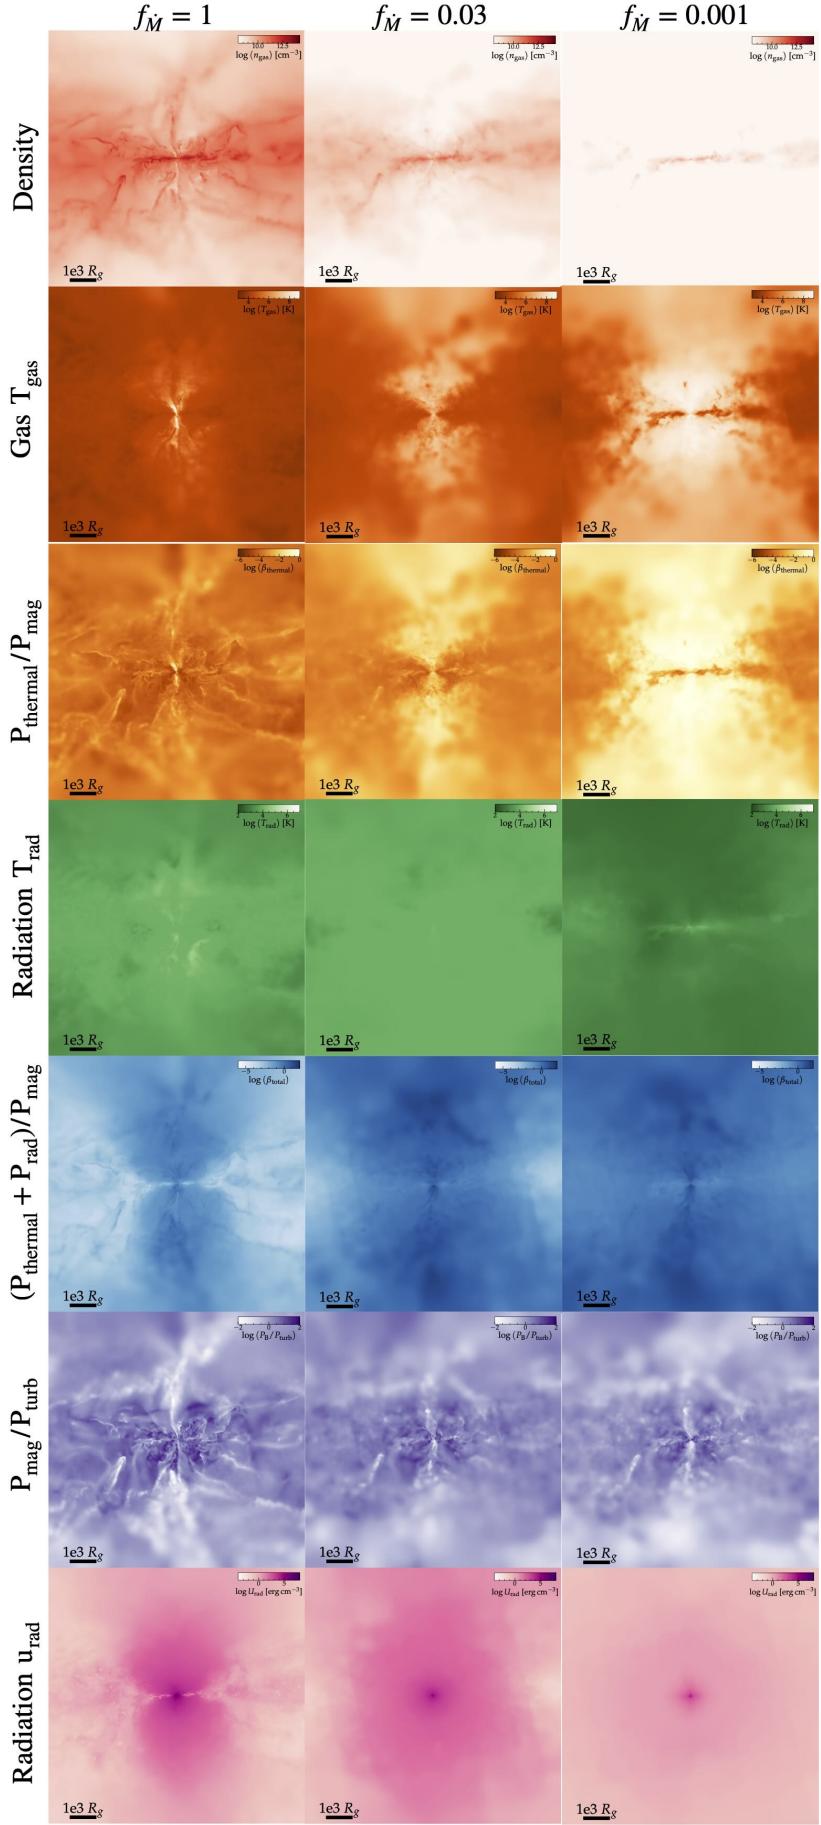


FIG. 21.— Edge-on images as Fig. 13, but comparing runs with different accretion rates $f_{\dot{M}} = (1, 0.03, 0.001)$ or $\dot{m} \sim (20, 0.3, 0.02)$ (left, middle, right) at fixed spatial scale. Mass and radiation density decrease with lower \dot{m} as expected, while β increases weakly, and M_A remains similar (see Fig. 16). Vertical temperature structure in T_{gas} and T_{rad} becomes more prominent, with a more distinct “sandwich” corona in T_{gas} and vertical T_{rad} stratification at lower \dot{m} .

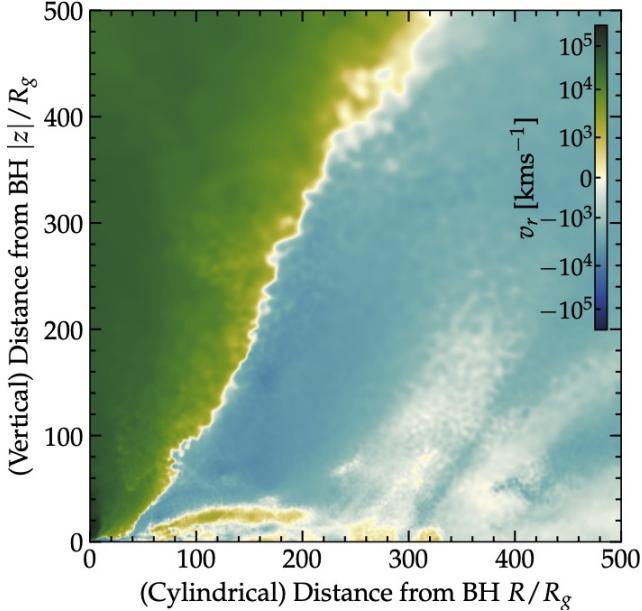


Fig. 22.— Illustration of the bipolar radiation-pressure driven outflows in the inner disk for our super-critical ($f_M = 1$) case. We plot $|z|$ versus R (out to $500 R_g$) to collapse to a single quadrant and show colors indicating the radial velocity (labeled). Most of the volume, especially beyond this radius, shows inflow onto and through the disk. But there is a clear outflow bicone as seen in Fig. 9, narrowing at larger r . The outflow velocities range from relativistic in the inner tens of R_g to $\gtrsim 10^4 \text{ km s}^{-1}$ at $\sim 500 R_g$ along the polar axis or $\sim 1000 \text{ km s}^{-1}$ along the edge of the cone/disk surface.

large-scale coherent, escaping flows.

The case of greatest interest here is our simulation with $\dot{m} \gg 1$ ($f_M = 1$), which does exhibit large-scale, coherent continuum radiation pressure driven outflows at the times shown (e.g. Figs. 9, 18, 22). In Fig. 11 we see that radiative acceleration is significantly weaker than gravity in the disk midplane. This is expected – the disk itself is optically thick so is externally self-shielding (Figs. 4, 12, 16), and is too heavy to be accelerated into radiation-pressure-driven outflow (Fig. 9) given the saturation of the radiation pressure from within the disk (§ 4.1) producing locally sub-Eddington fluxes from each annulus (Fig. 11) in the outer disk. As seen in other supercritical simulations (Jiang & Dai 2024, and references therein), the radiative acceleration is able to act most strongly in the less-dense, less Thompson-thick material in the polar direction, so the outflows driven tend to be strongly bipolar. The outflows feature velocities broadly comparable to the circular velocity at the corresponding radii, from $\sim 10^3$ to $\gtrsim 10^5 \text{ km s}^{-1}$ at $\lesssim 100 R_g$ (Fig. 22), with a strong internal shear (leading the large radial Reynolds term $\langle \delta v_r \delta v_r \rangle$ in Figs. 8 & 17), as expected (e.g. Murray et al. 1995). And the total mass outflow rate (Figs. 5, 16) is generally order-of-magnitude comparable to, but a factor of a couple to a few smaller than, the inflow rate (similar to other recent simulations with very different numerical methods; see Kurosawa & Proga 2009; Kuiper et al. 2012; Jiang et al. 2019; Toyouchi et al. 2021; Shi et al. 2024a,b; Jiang & Dai 2024). However the outflow rates can become larger in transient flare or outburst events (§ 3.5).

This means such outflows could be very significant in an “absolute” sense: carrying $\dot{M}_w > 1 M_\odot \text{ yr}^{-1}$ in mass, a momentum-flux $\dot{p}_w \sim M_w v_w$ comparable to $\sim L_{\text{bol}}/c$ (given the numbers above), and an at-launch energy flux

$\dot{E}_w \sim 0.5 \dot{M}_w v_w^2 \sim 0.01 - 0.05 L_{\text{bol}}$, comparable to values often invoked in AGN feedback models for self-regulation of both SMBH masses and galaxy properties (Silk & Rees 1998; Murray et al. 2005; Begelman & Nath 2005; Croton et al. 2006; Hopkins et al. 2006b,a; Ciotti et al. 2009). But they do not radically alter the structure of the disk itself, except perhaps in the most extreme burst/flare events (see § 3.5 and Fig. 15).

4.8. Properties of the Broad-Line Region

In Hopkins (2024), we predicted that the basic properties of the BLR should emerge naturally from the surface layers of geometrically-thick, self-illuminated magnetically dominated disks. We will study this in more detail in the future, but it is worth briefly commenting on the global properties of the gas in our simulations at the radii expected for the BLR at the different luminosities simulated ($\sim 1 - 100$ light-days, increasing with \dot{m} , e.g. Kaspi et al. 2005). We see the gas at these radii (1) has typical velocities $\sim 10^3 - 10^4 \text{ km s}^{-1}$; (2) exhibits covering factors $H/R \sim 0.1 - 0.3$; (3) is partially-ionized/partially-atomic; (4) has temperatures from few thousand to $\sim 10^5 \text{ K}$; and (5) densities $\sim 10^9 - 10^{13} \text{ cm}^{-3}$, with gas-mass-weighted mean $\sim 10^{10} - 10^{11} \text{ cm}^{-3}$; (6) exhibits clear clumping/substructure at least down to scales $< 3 \times 10^{12} \text{ cm}$ (limited by our spatial resolution; the minimum predicted Sobolev length given the resolved disk properties from the expressions in Hopkins 2024 ranges from $\sim 10^{10} - 10^{13} \text{ cm}$); (7) has $|\mathbf{B}| \ll G$ at the largest radii $\gtrsim 0.1 \text{ pc}$ (well below maser Zeeman upper limits where they exist; see Modjaz et al. 2005; Henkel et al. 2005; Vlemmings et al. 2007); (8) has explicitly-calculated absorbed/reprocessed light fractions $\gtrsim 10\%$; (9) show modest optical depths from disk surface, implying line re-emission can escape from illuminated “layer”; (10) has typical ionization parameters (combining the plotted densities and flux values) of $\sim 0.001 - 1$ (mostly varying with the large local density fluctuations in Fig. 5); (11) has a total gas mass $\sim 10 - 40 M_\odot$; and naturally (12) exhibits a “thick disk” geometry and kinematics. All of these properties agree remarkably well with what is observationally measured for the BLR (Krolik et al. 1981; Arav et al. 1998; Krolik 1999; Kaspi et al. 2005; Laor et al. 2006; Peterson 2006; Gravity Collaboration et al. 2018; GRAVITY Collaboration et al. 2020; Devereux 2021; GRAVITY Collaboration et al. 2021, 2024). Clearly, a more quantitative analysis and comparison to observations is merited: but this requires post-processing the simulations with radiative line transfer and detailed metal ionization state calculations, and therefore is deferred to future work.

4.9. Radiation-Thermal States & Deviations from Photon-Trapped LTE

Given our treatment of radiation-thermochemistry which allows for non-steady-state, non-LTE, multi-group, independent T_{gas} , T_{rad} , T_{dust} temperatures and non-thermal processes (§ A), as well as incomplete gas-radiation coupling, it is worth examining the different pressures and deviations from LTE or tight coupling in the inner disk. This is a particularly interesting region, because historical radiation-MHD treatments (often motivated by SS73-type models, with orders-of-magnitude larger disk optical depths and densities) have often assumed single-radiation-group, thermal-only/blackbody, strictly LTE, perfect gas-radiation coupling (in addition to neglecting atomic/molecular/dust processes).

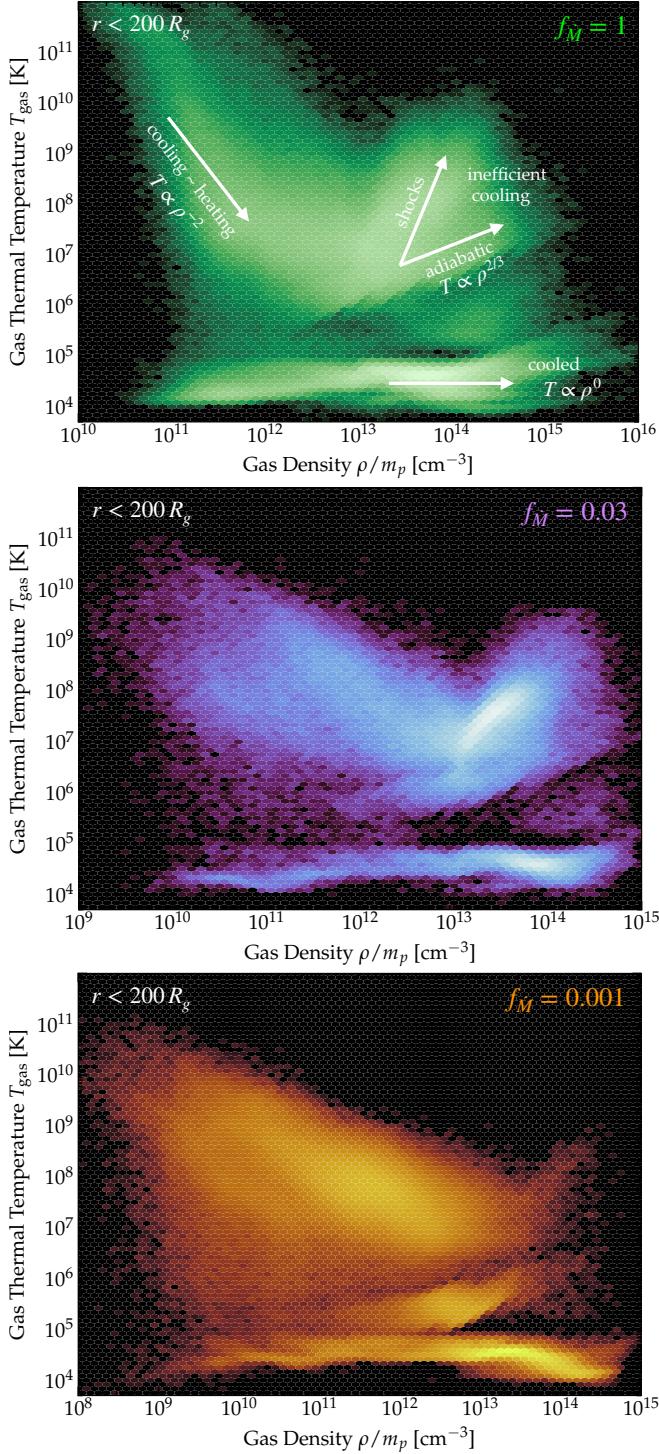


FIG. 23.— Gas temperature T_{gas} -density ρ phase diagram in the inner disk ($r < 200 R_g$), for our three simulations with different \dot{m} at the same time (labeled). Even in this region, complicated phase structure appears. We broadly label different behaviors. Hot and low density gas is locally optically-thin with optically-thin radiative cooling ($\propto n^2 \Lambda$) balancing external heating from e.g. turbulence, shocks, and reconnection ($\propto \rho \delta v_t^2$), giving an inverse $T_{\text{gas}} \propto \rho^{-2}$ -like scaling. Efficiently-cooled gas sits near temperatures regulated by local collisional plus photo-ionization equilibrium, giving near isothermal $T_{\text{gas}} \propto \rho^0$ behavior. Sufficiently hot, high-density ($n \gtrsim 10^{13} \text{ cm}^{-3}$) cannot cool rapidly and so is heated adiabatically ($T_{\text{gas}} \propto \rho^{2/3}$) or more rapidly still in shocks. All runs exhibit similar behavior despite different \dot{m} , though the density scales shifts. At larger radii, multi-phase structure becomes even more pronounced, and atomic/molecular phases eventually appear at $T_{\text{gas}} \lesssim 10^4 \text{ K}$.

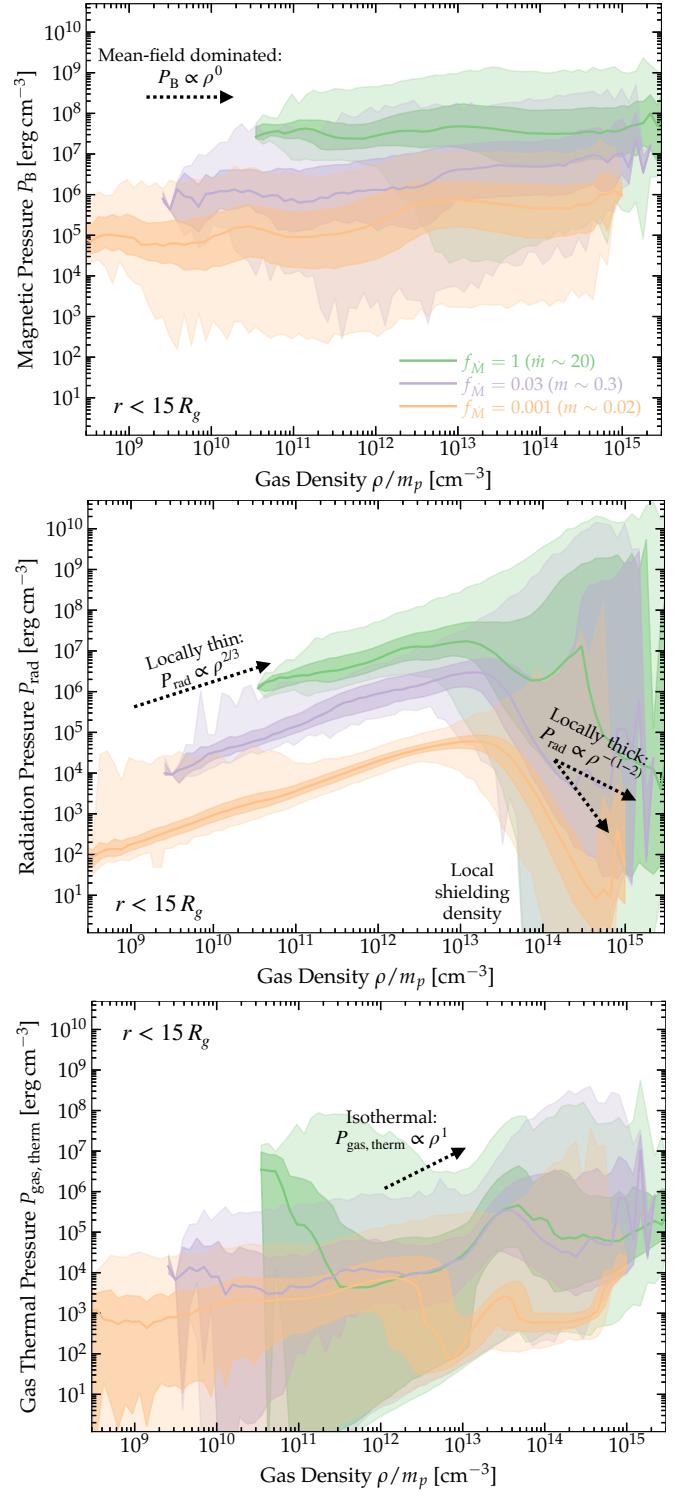


FIG. 24.— Pressures P versus gas density ρ in the innermost disk ($r < 15 R_g$) for our three simulations with different \dot{m} at the same time. We show the volume-weighted median (line), $\pm 1 \sigma$ (dark shaded), and $\pm 3 \sigma$ (light shaded) P at each ρ . Top: Magnetic pressure P_B . The strong mean toroidal field dominates so this depends weakly on ρ at a given R , though trans-Alfvénic turbulence generates significant fluctuations. Middle: Effective radiation pressure $P_{\text{rad, eff}}$ (§ 4.9). At a given R , we see a soft $P_{\text{rad}} \propto \rho^{2/3}$ approximately at lower densities, tracing locally-optically-thin gas where not all of u_{rad} can couple strongly, which turns at a critical density and then falls rapidly as $\propto \rho^{-1} - \rho^{-2}$, tracing dense structures at the thermal scale-length which are strongly shielded from the external radiation. Bottom: Gas thermal pressure $P_{\text{gas, therm}}$. Very crudely the median follows an isothermal-like track but with large deviations from multi-phase structure (§ 23).

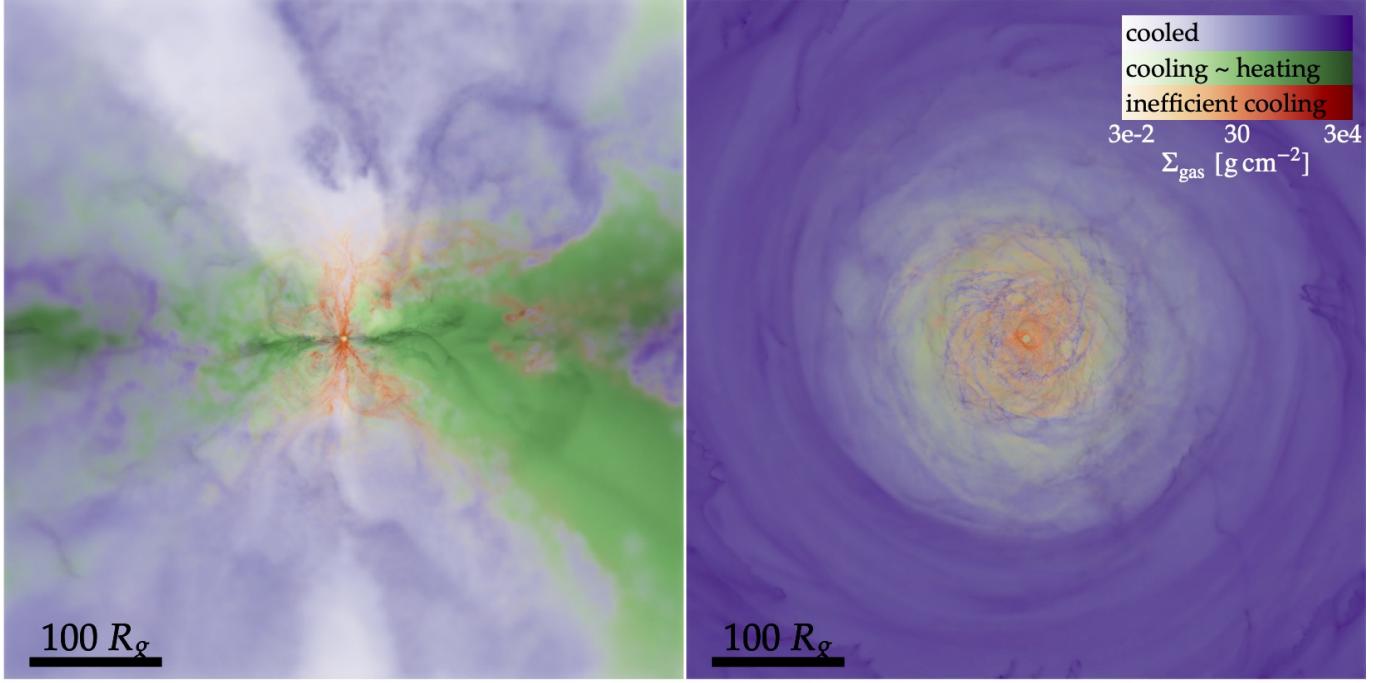


FIG. 25.— Edge-on and face-on projections (as Fig. 13) of gas surface density Σ_{gas} of the high- \dot{m} ($f_{\dot{M}} = 1$) snapshot from Fig. 23, within the same radii, highlighting the spatial distribution of gas in the different phases labeled in Fig. 23. Specifically gas which is quasi-isothermal at $T < 10^6$ K (“cooled”), or warmer gas at either low densities ($\lesssim 10^{12.5} \text{ cm}^{-3}$, “cooling \sim heating”) or higher densities ($\gtrsim 10^{12.5} \text{ cm}^{-3}$, “inefficient cooling”). At larger radii more of the gas is in cooler phases (as expected), but all three co-exist at similar radii, with some cold/cooled dense structures in the midplane, others cooling but hotter, and the highest surface-density (hence highest optical depth) structures in the jet/outflow shock regions showing less efficient cooling.

We will see large deviations from these behaviors, as predicted in e.g. Hopkins (2024), because these disks are only *globally* modestly optically thick to self-absorption (Fig. 11). This means large-scale modified blackbody emission is expected, but that *local* deviations from equilibrium and tight coupling, e.g. in density structures which are much smaller than the disk (which we see are ubiquitous, § 3) and therefore have much lower local/self optical depths, are also expected. Note that we focus on the more steady-state disk behaviors here, but the deviations from LTE and thermal radiation can become even larger in transient flare/outburst-type events (§ 3.5).

4.9.1. Gas Thermal Phase Structure

Fig. 23 shows the thermal gas phase structure in the inner disk (the multi-phase structure in the outer disk at $> 300 R_g$ is discussed extensively in Paper I, Paper II, and Paper III). The gas is clearly not single-phase, with at least three qualitatively different behaviors of a significant fraction of the mass/volume. At low densities and high ($\gtrsim 10^6$ K) temperatures, the gas is locally optically-thin, with temperature set by approximate equilibrium between volumetric optically thin cooling (mostly free-free, so $\dot{u}_{\text{therm}} \sim n^2 \Lambda \propto \rho^2 T_{\text{gas}}^{1/2}$) balancing heating by shocks/turbulence/magnetic reconnection (all similar, given the trans-Alfvénic turbulence, with $\dot{u} \sim \rho \delta v_{\text{turb}}^3 / L_{\text{turb}} \sim \rho v_A^2 \Omega$), so we anticipate something like $T_{\text{gas}} \propto \rho^{-2}$, broadly similar to what is seen. Gas which has cooled or is cooling even more efficiently sits at the collisional plus photo-ionization equilibrium at $T_{\text{gas}} \sim 10^4 - 10^5$ K as expected, and is nearly isothermal. This validates our expectation from analytic models (Hopkins 2024) that $t_{\text{cool}} \ll \Omega^{-1}$, even in the optically thin gas, throughout almost all of the disk (except some hot coronal gas). We have also verified this di-

rectly by checking the cooling rates in-code. Hotter, denser gas becomes optically-thick to its own cooling radiation *locally*, so has cooling times longer than its local dynamical/crossing times and so can be heated adiabatically ($T_{\text{gas}} \propto \rho^{2/3}$, as it is effectively monatomic here being fully-ionized) or in non-radiative shocks (which give density enhancement of a fixed factor of a few, with large temperature enhancement $\sim M_s^2$, so rise more steeply).

We can obtain a simple estimate for the critical density $n_{\text{crit}} \sim 10^{13} m_p \text{ cm}^{-3}$ of the transition in Fig. 23, if we ask when dense structures or shocks at the thermal scale length $H_{\text{therm}} \sim c_s / \Omega$ will become opaque, $\tau_{\text{therm}} \sim \kappa_{\text{es}} H_{\text{therm}} n_{\text{crit}} \gtrsim 1$, giving $n_{\text{crit}} \sim 10^{13} \text{ cm}^{-3} (R/10 R_g)^{-3/2} (x_e M_{\text{BH}}/10^7 M_{\odot})^{-1} (T_{\text{gas}}/10^7 \text{ K})^{-1/2}$.

4.9.2. Pressures versus Density

Fig. 24 compares different pressures versus density in the inner disk more broadly-defined. As expected and seen in Paper II, given that we see the mean field generally dominates $|\mathbf{B}|$, the magnetic pressure as usually defined $P_B \equiv |\mathbf{B}|^2/8\pi$ depends only weakly on ρ at a given radius (almost all the weak trend in P_B versus ρ in Fig. 24 comes from the dependence of both on R within the range plotted). But trans-Alfvénic turbulence does produce significant fluctuations. Magnetic pressure is not isotropic: this means as discussed in the main text, Paper II, and Hopkins et al. (2024a), that gas will generally “feel” a stiff magnetic equation-of-state in compression in the vertical or radial (perpendicular) directions, but soft in the azimuthal (parallel) direction, though the latter compressions are strongly sheared-out.

The thermal pressure in Fig. 24, $P_{\text{gas, thermal}} \equiv (\rho/\mu m_p) k_B T_{\text{gas}}$, has a vaguely isothermal “floor” as expected from Fig. 23 but reflects the same complicated multiphase

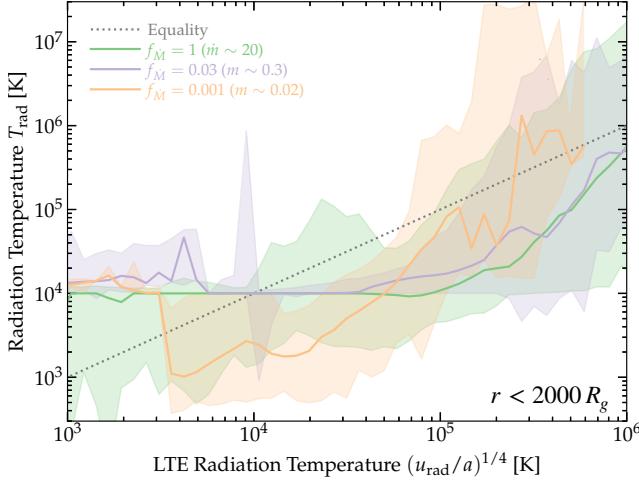


FIG. 26.— Comparison of the explicitly-evolved simulation radiation temperature T_{rad} in each cell in the inner disk at $r < 2000 R_g$ versus the expected radiation temperature in LTE $T_{\text{rad, LTE}} \equiv (u_{\text{rad}}/a)^{1/4}$, for our three different \dot{m} simulations with volume-weighted median (line) and $\pm 3 \sigma$ range (shaded) at each $T_{\text{rad, LTE}}$. While order-of-magnitude similar, the values of T_{rad} at a given u_{rad} can vary by ~ 2 dex and the systematic offset can be as large as ~ 1 dex. Even in the inner disk, it therefore important to evolve T_{rad} explicitly.

structure.

Less trivial is the radiation pressure. Here we define $P_{\text{rad, eff}, i} \equiv [\tau_{\text{local}, i}/(1 + \tau_{\text{local}, i})] u_{\text{rad}, i}/3$ for each cell, where $\tau_{\text{local}, i} = \langle \kappa_i \rangle_v \rho_i / |\nabla \rho_i|$ reflects the local flux-weighted mean opacity and density gradient scale-length. The usual convention in the accretion disk literature is to simply take $P_{\text{rad}} \equiv u_{\text{rad}}/3$, and this used in e.g. the standard definition of β_{tot} in the text. But as noted there, this can be very misleading in some limits, like when the gas is locally optically-thin, and there can be arbitrarily small radiation forces/accelerations on the gas at any given P_{rad} . More rigorously, we can compare the forces directly as in Fig. 11, where we see P_{rad} only dominates the actual *acceleration* for gas in the inner disk in the highly polar regions for the highest \dot{m} . But the definition in Fig. 24 gives us a better proxy for the “effective” radiation pressure which can actually act on a parcel of gas. Notably, we see a transition at a similar n_{crit} discussed in § 4.9.1. Below this density, we expect local structures to be *locally* optically thin, so the full radiation energy density cannot effectively act as a pressure. Even if u_{rad} is homogeneous, we expect lower $P_{\text{rad, eff}}$ at lower densities owing to the weaker coupling, and see this, with effective $P_{\text{rad, eff}} \propto \rho^{2/3}$. Above the critical density, we see a huge spread in P_{rad} . Structures here are self-shielded against the ambient radiation field. Thus at the highest P_{rad} we see hot, non-cooling structures from Fig. 23, while at the lowest P_{rad} we see structures at the disk mid-plane/dense clumps which are cooling (often $\sim 10^4 - 10^5$ K dense gas from Fig. 23) but shielded from the collective radiation field. This is the natural result of the disk being *externally* illuminated over much of its size, with lower- P_{rad} center seen in Figs. 4, 12, 20.

There appears, therefore, to be no regime in which one can “safely” take the tight-coupling limit, which would predict $P_{\text{rad, eff}} \propto \rho^{4/3}$. This also means that one cannot with any meaning define an “effective single-fluid sound-speed” $c_{s, \text{eff}}$ (commonly taken to scale as $c_{s, \text{eff}} \approx \sqrt{\gamma_{\text{eff}} [P_{\text{rad}}/P_{\text{therm}}] (P_{\text{rad}} + P_{\text{therm}})/\rho}$ such that $\gamma_{\text{eff}} \rightarrow 5/3$ for $P_{\text{rad}} \ll P_{\text{therm}}$ and $\gamma_{\text{eff}} \rightarrow 4/3$ for $P_{\text{rad}} \gg P_{\text{therm}}$; see Mi-

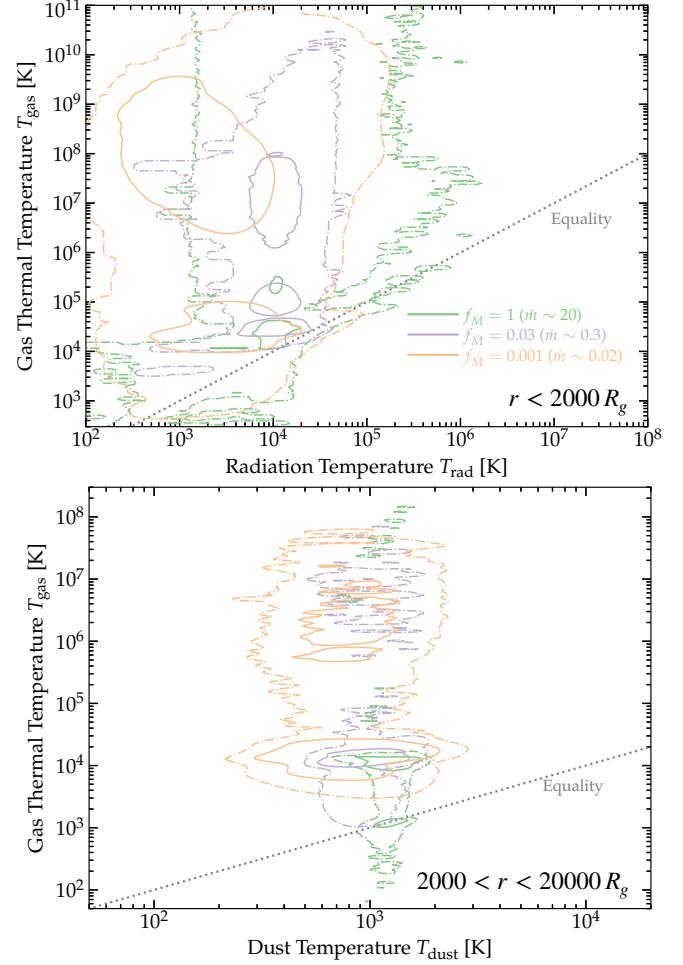


FIG. 27.— Comparison of different evolved temperatures in the inner disk. For each, we show contours corresponding to the volume-weighted 2D histogram at iso-density contours near-peak (solid) and enclosing $> 99\%$ of the volume (broken). Top: Gas thermal temperature T_{gas} versus T_{rad} . Bottom: T_{gas} versus dust temperature T_{dust} , only for cells with $> 1\%$ of their potential dust mass (not sublimated; this requires we look at larger radii $> 2000 R_g$ to see). The radiation, gas, and dust temperatures are highly NLTE, even in the central/inner disk at the highest \dot{m} , where the optical depths are maximized.

halas & Mihalas 1984), as these are only really valid in the infinite-optical-depth LTE and strong tight-coupling regime.

We can also verify that the qualitative behavior in Fig. 24 is similar, albeit with systematically lower pressures and densities, at different radii (making the same plot with the gas at e.g. $r = (15 - 30, 50 - 150, 150 - 300) R_g$). The predicted n_{crit} does shift weakly with radius as expected from our simple argument in § 4.9.1. Outside this radius the radiation pressures are even weaker and gas more strongly multiphase (Paper II).

4.9.3. Temperatures versus LTE

Fig. 26 plots one diagnostic of LTE in the inner disk, specifically comparing the explicitly-evolved simulation radiation temperature T_{rad} , to the LTE expectation $T_{\text{rad, LTE}} \approx (u_{\text{rad}}/a)^{1/4}$. If the radiation is either not perfectly thermal/blackbody, or the system is not in strictly LTE (i.e. does not have absorption mean-free-path much smaller than other scale-lengths), these will not necessarily agree. Consistent the behaviors above and the disks being only modestly optically-thick, we see that while T_{rad} , on average, is not wildly different from the LTE expectation, the local deviations can be large

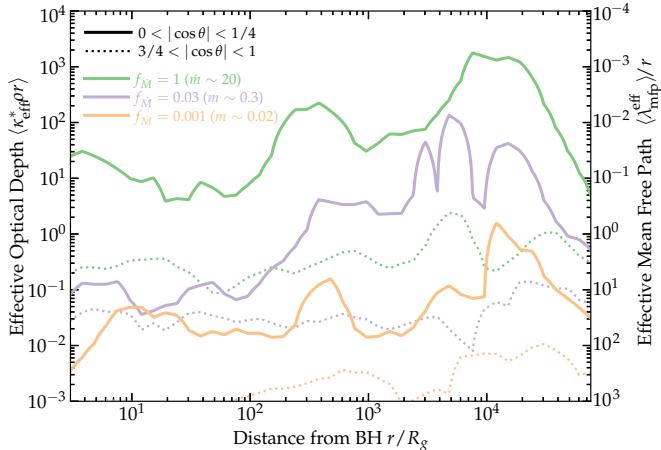


Fig. 28.— Characteristic local effective optical depth, $\kappa_{\text{eff}}^* \equiv \sqrt{\kappa_{\text{abs}}(\kappa_{\text{abs}} + \kappa_s)}$ in terms of absorption κ_{abs} and scattering κ_s opacities. As Fig. 11, we plot the median (volume-weighted) $\kappa_{\text{eff}}^* \rho r \sim r/\lambda_{\text{mfp}}^{\text{abs,eff}}$ in each annulus for cells within the given polar angle range (labeled). Note that calculating the integrated $\tau_{\text{eff}}^* \equiv \int \kappa_{\text{eff}}^* \rho d\ell$ for photon trajectories from e.g. the midplane at r to escape can give factor of $\gtrsim 10 - 30$ larger τ_{eff}^* in the inner disk; the plotted number is better thought of as the ratio r to the median effective absorption mean-free path $\lambda_{\text{mfp}}^{\text{abs,eff}}$ seen by photons instantaneously at some r and θ . This is always smaller than the total optical depth ($\kappa_{\text{abs}} < \kappa_s$), and smaller than the global optical depth to escape, but shows that much of the volume in the inner disk is *locally* optically-thin to absorption in local density structures with scale-lengths $\lesssim H \ll r$.

(up to $\gtrsim 1$ dex in either direction).

Fig. 27 further compares T_{gas} to T_{rad} , as well as T_{dust} (where dust is not completely sublimated) to T_{gas} . There is little or no correlation between the three, though T_{dust} correlates better with T_{rad} at much larger radii outside the sublimation radius (where dust reprocesses most of the radiation; Paper I) and T_{gas} does exhibit a lower envelope around T_{rad} (it is rare to have $T_{\text{gas}} \ll T_{\text{rad}}$ in the inner disk, though this can be true in the outer disk; see Paper II).

4.9.4. Opacities and Criteria for LTE & Tight-Coupling

Fig. 28 plots a simple proxy for the characteristic local optical depths along different polar angles within the simulation, as Figs. 5 & 16, but instead of the total opacity $\kappa_{\text{tot}} \equiv \kappa_s + \kappa_{\text{abs}}$ (the sum of the scattering and absorption opacities), we plot the effective absorption opacity $\kappa_{\text{eff}}^* \equiv \sqrt{\kappa_{\text{abs}} \kappa_{\text{tot}}}$. Almost everywhere in the inner disk, $\kappa_{\text{abs}} \ll \kappa_s \sim \kappa_{\text{tot}}$, typically by factors of $\sim 100 - 1000$, so $\kappa_{\text{eff}} \sim (0.03 - 0.1) \kappa_{\text{tot}}$ (compare Fig. 16). In the outer disk (especially in the midplane at $\gtrsim 3000 R_g$), dust survives, with (given the warm radiation temperatures) $\kappa_{\text{abs}} \sim \kappa_s$, so the difference is minimized. Note that we plot a median local $\kappa_{\text{eff}}^* \rho r \sim r/\lambda_{\text{mfp}}^{\text{abs,eff}}$ (roughly, the ratio of r to the effective absorption mean-free-path $\lambda_{\text{mfp}}^{\text{abs,eff}}$) in different annuli – this can be smaller by as much as an order of magnitude compared to what one often refers to as the “optical depth of the disk” obtained by integrating $\int \kappa_{\text{eff}}^* \rho d\ell$ along a hypothetical photon trajectory from the midplane to infinity, because we are not interested here in the *global* optical depth and reprocessing in the disk and atmosphere, but rather in the *local* mean free-paths $\lambda_{\text{mfp}}^{\text{abs,eff}}$ of the photons. Essentially, this gives $r/\lambda_{\text{mfp}}^{\text{abs,eff}}$ the inverse of the median absorption mean-free-path length in units of r at the given r and polar angle. Even in the midplane at our highest accretion

rates/surface densities ($f_{\dot{M}} = 1$), $\langle \kappa_{\text{eff}}^* \rho r \rangle \sim r/\lambda_{\text{mfp}}^{\text{abs,eff}}$ is modest (~ 10), i.e. the typical effective absorption mean-free-path $\langle \lambda_{\text{mfp}}^{\text{abs,eff}} \rangle \sim 1/\langle \kappa_{\text{eff}}^* \rho \rangle \sim 0.1 r \sim (0.3 - 1) H$. For lower accretion rates $\dot{m} \sim 0.02 - 0.3$ ($f_{\dot{M}} = 0.001 - 0.03$) $\langle \kappa_{\text{eff}}^* \rho r \rangle$ can be below unity, i.e. $\langle \lambda_{\text{mfp}}^{\text{abs,eff}} \rangle > r$ interior to $\lesssim 200 R_g$, which is equivalent to saying that most of the gas is *locally* optically-thin to absorption over the typical local density gradient scale lengths.

For comparison, consider the requirements for local photon-trapping/tight-coupling and strict LTE. “Full” local trapping, as required for “effective equation of state” or “effective sound speed” or single-fluid gas-radiation models (as opposed to the much less stringent global trapping required for advection of radiation into the BH), requires the total/scattering mean-free-path $\lambda_s \sim 1/\kappa_s \rho$ be much smaller than the smallest density gradient length scales ℓ , but also that the photon escape/diffusion time $\sim \ell/(c/\tau(\ell)) \sim \ell^2 \kappa_s \rho / c$ be much longer than the characteristic evolution times on ℓ , given by e.g. turbulent crossing times $\ell/v_{\text{turb}}(\ell)$. As noted in the text and derived in Hopkins (2024), for a supersonically-turbulent disk the ℓ of greatest interest is the sonic/shock/Sobolev length (all similar) $\ell_{\text{sonic}} \sim (c_s/v_{\text{turb}}) H_{\text{therm}} \sim H/\mathcal{M}_s^2$ (where c_s and \mathcal{M}_s are defined with respect to the gas thermal sound speed). So the timescale criterion becomes $\kappa_s \rho R \gg (c/v_{\text{turb}}[\ell])(H/R)^{-1} \mathcal{M}_s^2 \gtrsim 10^5 - 10^8$ (i.e. $\lambda_{\text{mfp}}^{\text{abs,eff}} \lesssim 10^{-8} - 10^{-5}$), much larger than we see (Fig. 11). The criterion for strict LTE is dimensionally similar but even more demanding because it requires reprocessing, so we must exceed the same threshold with the effective opacity κ_{eff}^* instead of κ_s . So unless the disks were orders-of-magnitude more dense, we should expect to find (as we do) that local tight-trapping and LTE are not good local approximations.

4.10. Sensitivity to Physics

In the inner disks ($\ll 1000 R_g$), there are many physical processes formally included in our simulations which can be safely neglected, and a few which are of key importance. We review these briefly here.

4.10.1. Physics Which Can Be Safely Neglected

There is zero star formation at $\lesssim 0.05$ pc in our longer-duration Paper II simulations and these simulations, and negligible stellar mass (from migration of stellar orbits) in the inner disk, so simply removing all star formation and stellar evolution and stellar “feedback” physics yields identical results at the radii of interest here. Dynamical self-enrichment/pollution of metals is therefore also negligible. Related, local self-gravity of the disk is weak as the inner disk mass is $\ll 10^{-6} M_{\text{BH}}$ and Toomre Q ranges from $Q \sim 10^5 - 10^{10}$ (though it remains possible that inward-propagating, swing-amplified self-gravity modes from large radii could be important; see Hopkins 2010; Lin 2015; Duffell et al. 2024).

Also per Paper II & Paper III, even the atomic disk is sufficiently well-ionized that non-ideal (Ohmic resistivity, Hall drift, ambipolar drift) effects are smaller than turbulent or numerical resistivities and reconnection rates by factors of $\sim 10^{15}$. Likewise the fluid is extremely well into the collisional regime with electron collision frequencies $\sim 10^{-10} - 10^{-6} \Omega$, so collisionless or kinetic effects and microphysical Spitzer/Braginskii/atomic/molecular conduction and viscosity are negligible. Similarly in the *disk*, the temperatures

are generally well below those where two-temperature effects (treated in a simplified manner in the simulations where relevant) are significant, though this is not true for all the gas in the diffuse disk atmosphere/corona (Fig. 5). And though we include a detailed non-equilibrium molecular chemical network, there is relatively little molecular gas at $< 300 R_g$ except at the lowest \dot{m} (though atomic and molecular gas are certainly important in the outer disk). The external cosmic ray background is also weak on these scales (for both bulk disk support, and direct cosmic ray heating/cooling processes, though it may help non-trivially in supporting the relatively large ionization fractions), with CR energy densities falling below the plotted range in Fig. 5. Dust is largely sublimated at the interior radii $\sim 30 - 300 R_g$, even in the self-shielded midplane of the disk, though it remains possible that with a more sophisticated model for dust re-formation some of the colder-phase gas we see could experience grain re-growth. So cosmic rays, dust, and detailed molecular thermo-chemistry are not directly dynamically important if one wishes to *only* focus on small scales $\ll 300 R_g$.

On most of the scales we resolve, relativistic effects are generally small. For example (§ B) it makes little difference if we treat the SMBH with a Keplerian or Paczyński-Wiita potential at most radii, though it does become significant around the ISCO. But such effects could be important indirectly, e.g. Lens-Thirring effects tilting the inner disk, and will grow rapidly in importance at smaller radii. And they could be non-linearly important, as e.g. GR effects could help launch jets, which could in turn carve a channel for radiation to more efficiently escape from the inner disk (see Kaaz et al. 2024).

Numerically, we refer to § B to illustrate that numerical choices such as the details of our numerical refinement scheme, resolution, exact treatment of different radiation bands and opacities/cooling rates, and when precisely we select to “zoom in” from the parent simulation, do not qualitatively change our conclusions.

4.10.2. Most Important Physics

The most important physics “ingredients” then (given initial and boundary conditions) for our focus here are MHD and radiation-thermochemistry in a BH potential. The MHD (being well-approximated by ideal MHD) and gravity (from the BH) are straightforward. But the radiation-thermochemistry undergoes several transitions over the range of radii we evolve: between optically thin and thick, radiation pressure-dominated and not, thermalized and non-thermalized radiation, atomic and ionized gas, scattering and absorption opacity-dominated, IR and optical/UV emission-dominated. It is therefore important to utilize radiation-thermochemistry methods which can flexibly treat these different limits.

Removing any radiation flux (in the Lagrangian frame), i.e. “locking” radiation to gas leads to the disk heating adiabatically and then “exploding” or “pulsing” (launching a radial “bounce”) on of order one dynamical time, as expected – this simply confirms that escaping radiation can be an important energy “sink.” Conversely, simply deleting all cooling radiation (i.e. assuming completely optically-thin escape of photons) so that cooling remains but there is zero radiation pressure, as shown in Figs. 29 & 30 leads to a colder, neutral atomic+molecular disk at large radii, which can become somewhat thinner/denser and less strongly-turbulent in the inner regions, and naturally never launches radiation-pressure driven outflows even for $\dot{m} \gg 1$. So radiation plays some significant roles here, but the effects are somewhat modest,

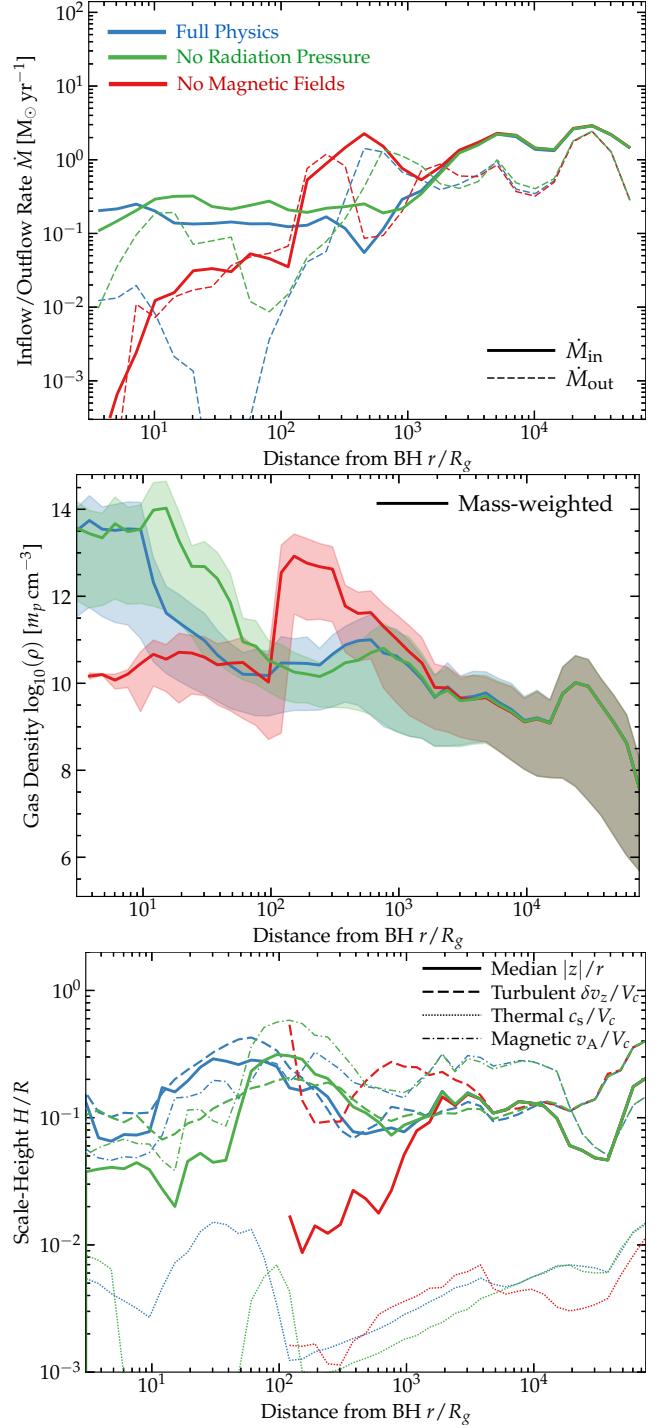


FIG. 29.— Radial profiles of inflow/outflow rates \dot{M} (top), densities (middle), and scale-heights (bottom) as Figs. 5 & 16, but comparing our default/“full-physics” $\dot{m} \sim 0.3$ ($f_{\dot{M}} = 0.03$) simulation to two restarts (all at the same time since restart). In one, we remove radiation-hydrodynamics and assume optically-thin cooling (all cooling radiation is simply lost), so there is no radiation pressure. In another, we retain full radiation but remove magnetic fields. Both have significant effects (§ 4.10.2) in the inner disk (the larger radii being identical is simply an artifact of the short run-time of the simulations, showing the effects propagate inside-out with time, as expected). Removing radiation pressure leads to a somewhat thinner, less turbulent, cooler inner disk, with longer accretion times, larger in-plane outflow/radial sloshing motions, and higher densities at the same \dot{M} . Removing magnetic fields leads to collapse of the inner disk: the inflow rates drop by factors ~ 100 and are equal to outflows from weak eccentric/spiral/sloshing modes, an expending cavity forms, and the scale-height drops to nearly the resolution limit (we do not plot the “no magnetic field” case in H/R interior to the cavity, because the density scale-height becomes unresolved).

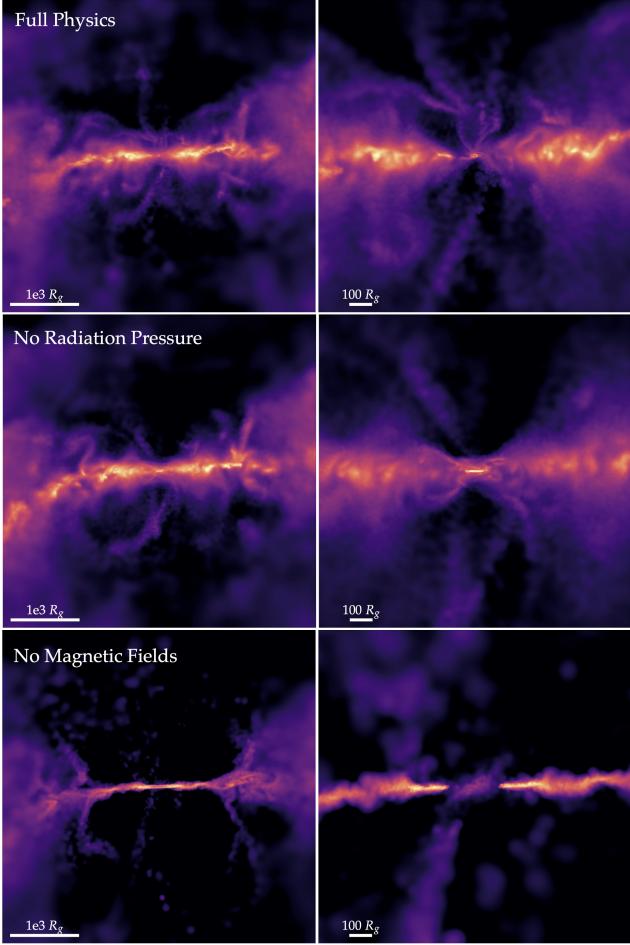


FIG. 30.— Images of the simulations from Fig. 29, comparing our default/“full-physics” $\dot{m} \sim 0.3$ run (top) to restart at the same time with no radiation pressure (middle) or no magnetic fields (bottom), at two different scales (left and right). Without radiation pressure we can see a slightly thinner very inner disk at $\ll 100 R_g$, the lack of the vertical (radiation-pressure driven) outflows from the very central disk, and a slightly smoother (i.e. less-turbulent/clumpy) disk on scales $\lesssim 1000 R_g$. Without magnetic fields, the collapse to a much thinner disk and presence of a central hole/cavity (out to $\sim 100 R_g$) owing to the *decretion* disk behavior are obvious.

even in regions where a naive comparison of P_{rad} , P_{thermal} , P_B would imply radiation pressure is dominant. This is discussed above in § 4.9, where we show that despite the disk being globally optically-thick and advecting photons into the BH, “tight” or “local” photon trapping (i.e. assuming that photon diffusion/escape times are much shorter than any other timescales in the problem even in small density sub-structures within the disk), as required for modeling the disk with some simply $\gamma = 4/3$ “effective” equation of state or “effective” gas+radiation single-fluid sound speed, is not a good approximation. This cautions against such approximations, as well as showing that even more demanding approximations (like assuming the radiation and opacities/excitation states of different species all obey LTE) are not accurate in these simulations.

Figs. 29 & 30 show that removing magnetic fields has dramatic, leading-order effects, similar to those seen on larger scales in Paper II. In order to start our simulation from a self-consistent zoom-in lacking magnetic fields, we could take the simulation without magnetic fields (zooming in from Mpc to $\sim 300 R_g$ scales) presented in Paper II, and simply move the inner boundary inwards to $\sim 30 R_g$ (keeping its resolution

$\Delta m \approx 5 \times 10^{-8} M_\odot$). As shown in Paper II, if we do this, we find it leads to catastrophic fragmentation of the outer accretion disk, with SFRs higher by a factor of ~ 1000 at $\sim 0.1 - 10$ pc and a factor of $\sim 10^5 - 10^6$ at $\sim 0.01 - 0.1$ pc. The combination of this and orders-of-magnitude weaker stresses (weaker by a factor of $\sim 10^2 - 10^4$) means that outflow and SFRs are then much larger than inflow rates and the disk is starved, producing an inflow rate at least a factor of ~ 1000 times lower into $< 300 R_g$ (and indeed for most of the time, as shown in Paper II, the inner region has a net Reynolds stress producing outflow, hence forms an expanding “hole” in the inner region or *decretion* disk). So a “further zoom in” would simply be uninteresting as there is no gas at $< 300 R_g$. So instead in Figs. 29 & 30 we take one of our simulations here in which strong magnetic fields developed self-consistently, and restart after turning off MHD (i.e. “deleting the fields”). This is not self-consistent, but is a useful counterfactual: since turbulent, magnetic, and radiation pressures are order-of-magnitude comparable in the inner disk, i.e. $\beta_{\text{tot}} \sim 1$, one might imagine turbulence and radiation somehow sustaining the disk without fields. But we find that removing magnetic fields, two dramatic changes occur even in the inner disk: (1) the disk loses vertical support and collapses to a vertical density scale-height approximately equal to the resolution limit, and (2) the stresses driving accretion not only become much weaker, but even reverse sign, with $\dot{M}_{\text{out}} \approx \dot{M}_{\text{in}}$ at all radii where we evolve for much longer than the dynamical time. These comparable inflow-outflow rates are dominated by random turbulent and sloshing/eccentric motions, with a small net \dot{M}_{out} driving overall loss of mass inside of the inner disk while the time-averaged accretion rate actually reaching the ISCO/SMBH decreases by at least a factor of ~ 100 . Again, this is similar to the effects seen in our tests in Paper II, and we find again similar results with additional resolution tests without magnetic fields in Appendix D. So despite $P_{\text{rad}} \sim P_B$, the effects of removing radiation pressure and magnetic fields are very different.

4.10.3. Potentially Important Physics Neglected Here

While the list of physics included in these simulations is extensive, it is by no means exhaustive. As argued above, there are a couple pieces of physics that could be important in certain regimes. The most obvious is GR. This should only become important on near-horizon scales, but our lack of relativistic physics means (as discussed in § 4.7) that we cannot capture spin effects (important both for jets but also potentially misaligned disks and Lens-Thirring precession; see Kaaz et al. 2022), nor horizon-scale effects on the accretion of magnetic flux which could also influence jet-launching.

Also as noted in § 4.7, our multi-group radiation transport still lacks fine enough binning to properly treat line-driven winds. This could be critical for strong wind-launching from the surface of the disk in regions like the BLR. How important we expect this to be will be investigated in future work using more detailed Monte-Carlo radiation post-processing. We also adopt a simple model for dust sublimation, rather than explicitly following the evolution of dust grain populations (size evolution, destruction, condensation and coagulation), which could be important for the details of the location of the sublimation radii and for dust-radiation-pressure driven outflows from the inner edge of the torus (Soliman & Hopkins 2023).

And while we argued above that collisionless and two-temperature effects should be weak for the midplane accreting gas in the simulations here, it could definitely become impor-

tant for the hottest ($T_{\text{gas}} \gg 10^9$ K), most tenuous coronal gas around the ISCO. Here more detailed plasma physics treatments in idealized simulations (e.g. Bambic et al. 2024) are needed.

5. CONCLUSIONS

We extend multi-group radiation-magnetohydrodynamics simulations with self-gravity and detailed non-equilibrium and NLTE thermochemistry, originally run from \sim Mpc scales to follow gas forming a quasar accretion disk, down to ISCO scales. We also extend these to explore different inflow rates. The accretion disks studied involves accretion at $\sim 0.005 - 10 \text{ M}_\odot \text{ yr}^{-1}$ around a $\sim 10^7 \text{ M}_\odot$ SMBH, which for a reference radiative efficiency of $\epsilon_r^0 = 0.1$ translates to sub-to-supercritical accretion rates spanning $\dot{m} \sim 0.02 - 20$.

Most of the dynamical properties of the disks follow a reasonable extrapolation from larger radii: they maintain (thermal) $\beta \sim 10^{-5} - 10^{-2} \ll 1$, i.e. are strongly magnetically dominated, with a predominant toroidal mean magnetic field driven by flux freezing/advection (except in the innermost disk), with trans-Alfvénic and hence highly super-sonic and quasi-isotropic turbulent motions, and strong Maxwell and Reynolds stresses, in a thick, flared, multi-phase disk. All of these properties agree qualitatively with the analytic models presented in Paper III, designed to extrapolate from the BHROI at a few pc to near-horizon scales. The most important extension of Paper II is that we now resolve the transition radii where disks self-ionize, where most of the radiation is actually emitted, where radiation pressure can become important in the disks, and where (for the higher \dot{m}) a radiatively-efficient disk could become locally super-Eddington.

Here we show that even at $\dot{m} \gg 1$, supersonic turbulence makes the radiation transport globally advective, leading to a quasi-isothermal vertical structure with midplane temperature similar to the effective temperature (but with multiphase thermal structure within the disk), and $P_{\text{mag}} \gtrsim P_{\text{mag, turb}} \sim P_{\text{kin, turb}} \sim P_{\text{rad}} \gg P_{\text{thermal}}$, i.e. saturation of radiation and turbulent (kinetic+magnetic) pressures in equilibrium comparable to or a factor of a few smaller than the mean magnetic pressure (so the radiation+thermal $\beta_{\text{tot}} \equiv (P_{\text{rad}} + P_{\text{thermal}})/P_{\text{mag}} \sim 0.1 - 1$ in the inner disk regions). Closely related, at high \dot{m} the disk becomes radiatively inefficient at radii well interior to this transition, with a total emergent luminosity just a few times larger than Eddington. However we caution that any prediction for the “total” radiative efficiency at better than factor of a few accuracy may require GRMHD simulations that can more accurately extend to horizon scales, as much of the luminosity will be released on these scales or be carried back out by jets or other outflows (Kaaz et al. 2024), though it is also clear here that modeling the thick-to-thin transition and NLTE, non-tight-coupling limit of radiation is important for the radiation temperature, pressure, and escape. We also verify that these transitions towards advective transport modify the similarity solutions from Paper III for disk properties, as predicted in Hopkins (2024). We also show that despite $P_{\text{rad}} \sim P_{\text{B}}$ in the inner disk, radiation and magnetic fields act very differently on the gas, and removing magnetic fields causes much more dramatic effects on disk structure compared to removing radiation pressure. Large-scale coherent magneto-centrifugal winds do not emerge at any radii or accretion rates, but we show that this is expected given the large densities. However, “failed” magnetocentrifugal outflows or fountain flows could still be playing an important role driving turbulence and “puffing up”

the disk.

With the ability to resolve an order-unity fraction of the expected emergent bolometric luminosity, our simulations show that radiative properties of the disk are strongly influenced by these transitions as well as the geometrically thick, low-density, modest optical-depth structure generic to flux-frozen disks. The radiation flux escapes in a wide-angle bicone, but strongly self-irradiates the disk itself. All of the disk “surface layers” from radii $\sim 10 - 10^6 R_g$ are strongly illuminated, and an order-unity fraction (weakly decreasing at lower \dot{m}) of the entirely bolometric flux from the innermost disk is intercepted by and reprocessed or scattered in the outer disk. Even sightlines that naively would not intercept the disk (e.g. along more polar angles with $\cos \theta \gtrsim H/R$) can be scattered by the extended, accreting “atmosphere” of warm coronal gas at higher \dot{m} (which is expected, given that H/R is order-unity, there must be appreciable gas at $|z| \sim R$, and we see this in the slowly-declining vertical gas density profiles supported by even more slowly-declining magnetic pressure) back onto the disk. The thermal and radiation structure is complicated and multi-phase, and is not captured by LTE or single-fluid (tight gas-photon coupling) prescriptions. Continuum radiation pressure drives outflows in a relatively narrow, low-gas-density bicone from launching radii $\sim 10 - 100 R_g$ at the highest $\dot{m} \gtrsim 1$, without outflow rates order-of-magnitude comparable to (though usually somewhat below) the accretion rates in the inner disk.

These radiation behaviors will have important implications for observed quasar variability, reverberation mapping, microlensing sizes and their wavelength dependence, and the spectral shape of the quasar emission – as discussed in Hopkins (2024) it means the disk *itself* can easily provide the gas needed for historical “components” usually imagined to float above the (razor-thin) SS73-like accretion disk, like warm Comptonizing “skins”, the corona, the warm absorber, the broad line region, and dusty torus. Detailed predictions for the spectra at different wavelengths and their variability will require radiative transfer post-processing calculations using these simulations, as our simple few broad-band on-the-fly radiation-MHD algorithm is insufficient to make such predictions. But we can directly illustrate the salient scattering and reprocessing, and see, for example, its critical effect on the emergent quasar radiation field. For example, the innermost disk features a radiation temperature of almost $\sim 10^6$ K (roughly following $T_{\text{rad}} \propto R^{-3/4}$), but by the time we follow this radiation to radii $\gtrsim 2000 R_g$ where it is escaping in said bicone, the emergent effective temperature is $\sim 4 \times 10^4$ K, in excellent agreement with observed SEDs of quasars at similar luminosities. And we show that properties of the gas in the surface layers of the disk at radii corresponding to the BLR agree remarkably well with what is needed to explain the more detailed spectral properties observed.

The physics in our code, while designed to follow a large range of scales from $\sim 10^{14} - 10^{27}$ cm, have two key limitations that must be addressed at smaller radii. First, GR effects will become important, not just at the ISCO, but since (see Paper II) the disk here is likely strongly misaligned with the pre-existing BH spin, where Lense-Thirring precession becomes important. Second, our multi-group radiation methods, while including sophisticated optically thin and thick treatments for a wide range of chemical processes, do not include either the effects of hard X-ray over-ionization (of e.g. Fe) or line transfer, critical for understanding both the opacities and potential line-

driven winds launched by the inner disk. So the behavior of line-driven winds, the alignment of the central disk, and the launching of jets (presumably also near-horizon) remain important subjects for future work.

Support for PFH was provided by NSF Research Grants

1911233, 20009234, 2108318, NSF CAREER grant 1455342, NASA grants 80NSSC18K0562, HST-AR-15800. CAFG was supported by NSF through grants AST-2108230 and AST-2307327; by NASA through grants 21-ATP21-0036 and 23-ATP23-0008; and by STScI through grant JWST-AR-03252.001-A.

REFERENCES

Abramowicz M. A., Fragile P. C., 2013, *Living Reviews in Relativity*, **16**, 1
 Abramowicz M. A., Czerny B., Lasota J. P., Szuszkiewicz E., 1988, *ApJ*, **332**, 646
 Aller M. C., Richstone D. O., 2007, *ApJ*, **665**, 120
 Anderson J. M., Li Z.-Y., Krasnopolsky R., Blandford R. D., 2005, *ApJ*, **630**, 945
 Anglés-Alcázar D., et al., 2021, *ApJ*, **917**, 53
 Arav N., Barlow T. A., Laor A., Sargent W. L. W., Blandford R. D., 1998, *MNRAS*, **297**, 990
 Bacon R., Emsellem E., Combes F., Copin Y., Monnet G., Martin P., 2001, *A&A*, **371**, 409
 Bai X.-N., Stone J. M., 2012, *ApJ*, in press, arXiv:1210.6661
 Bai X.-N., Ye J., Goodman J., Yuan F., 2016, *ApJ*, **818**, 152
 Bambic C. J., Quataert E., Kunz M. W., Jiang Y.-F., 2024, *MNRAS*, **530**, 1812
 Begelman M. C., Nath B. B., 2005, *MNRAS*, **361**, 1387
 Benincasa S. M., Wadsley J., Couchman H. M. P., Keller B. W., 2016, *MNRAS*, in press, arXiv:1607.05795
 Benincasa S. M., et al., 2020, *MNRAS*, **497**, 3993
 Bisnovatyi-Kogan G. S., Ruzmaikin A. A., 1976, *Ap&SS*, **42**, 401
 Blackburne J. A., Pooley D., Rappaport S., Schechter P. L., 2011, *ApJ*, **729**, 34
 Blandford R. D., Payne D. G., 1982, *MNRAS*, **199**, 883
 Boller T., Brandt W. N., Fabian A. C., Fink H. H., 1997, *MNRAS*, **289**, 393
 Bonnerot C., Lu W., Hopkins P. F., 2021, *MNRAS*, **504**, 4885
 Cai Z.-Y., Wang J.-X., 2023, *Nature Astronomy*, **7**, 1506
 Choban C. R., Kereš D., Hopkins P. F., Sandstrom K. M., Hayward C. C., Faucher-Giguère C.-A., 2022, *MNRAS*, **514**, 4506
 Ciotti L., Ostriker J. P., Proga D., 2009, *ApJ*, **699**, 89
 Cochrane R. K., et al., 2023, *MNRAS*, **523**, 2409
 Cornachione M. A., Morgan C. W., 2020, *ApJ*, **895**, 93
 Crenshaw D. M., et al., 2000, *AJ*, **120**, 1731
 Croton D. J., et al., 2006, *MNRAS*, **365**, 11
 Czerny B., Nikołajuk M., Różańska A., Dumont A. M., Łoska Z., Zycki P. T., 2003, *A&A*, **412**, 317
 Dai X., Kochanek C. S., Chartas G., Kozłowski S., Morgan C. W., Garmire G., Agol E., 2010, *ApJ*, **709**, 278
 Davis S. W., Tchekhovskoy A., 2020, *ARA&A*, **58**, 407
 Dedner A., Kemm F., Kröner D., Munz C.-D., Schnitter T., Wesenberg M., 2002, *Journal of Computational Physics*, **175**, 645
 Deng H., Mayer L., Meru F., 2017, *ApJ*, **847**, 43
 Deng H., Mayer L., Helled R., 2021, *Nature Astronomy*, **5**, 440
 Devereux N., 2021, *MNRAS*, **500**, 786
 Dhang P., Bendre A. B., Subramanian K., 2024, *MNRAS*, **530**, 2778
 Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, **433**, 604
 Duffell P. C., et al., 2024, *ApJ*, **970**, 156
 Dunn J. P., et al., 2010, *ApJ*, **709**, 611
 Faucher-Giguère C.-A., 2020, *MNRAS*, **493**, 1614
 Faucher-Giguère C.-A., Quataert E., 2012, *MNRAS*, **425**, 605
 Faucher-Giguère C.-A., Quataert E., Murray N., 2012, *MNRAS*, **420**, 1347
 Ferrarese L., Merritt D., 2000, *ApJ*, **539**, L9
 Few C. G., Dobbs C., Pettitt A., Konstandin L., 2016, *MNRAS*, **460**, 4382
 Frank J., King A., Raine D. J., 2002, *Accretion Power in Astrophysics*: Third Edition, isbn 0521620538 edn. Cambridge, UK: Cambridge University Press, Cambridge, UK
 Frederick S., et al., 2021, *ApJ*, **920**, 56
 GRAVITY Collaboration et al., 2020, *A&A*, **643**, A154
 GRAVITY Collaboration et al., 2021, *A&A*, **648**, A117
 GRAVITY Collaboration et al., 2024, arXiv e-prints, p. arXiv:2401.07676
 Gaburov E., Johansen A., Levin Y., 2012, *ApJ*, **758**, 103
 Gandhi P. J., Wetzel A., Hopkins P. F., Shappee B. J., Wheeler C., Faucher-Giguère C.-A., 2022, *MNRAS*, **516**, 1941
 Gebhardt K., et al., 2000, *ApJ*, **539**, L13
 Glover S. C. O., Jappsen A. K., 2007, *ApJ*, **666**, 1
 Graham M. J., Djorgovski S. G., Drake A. J., Stern D., Mahabal A. A., Glikman E., Larson S., Christensen E., 2017, *MNRAS*, **470**, 4112
 Graham M. J., et al., 2020, *MNRAS*, **491**, 4925
 Gravity Collaboration et al., 2018, *Nature*, **563**, 657
 Grudić M. Y., 2021, *MNRAS*, **507**, 1064
 Grudić M. Y., Hopkins P. F., 2020, *MNRAS*, **495**, 4306
 Grudić M. Y., Hopkins P. F., Lee E. J., Murray N., Faucher-Giguère C.-A., Johnson L. C., 2019, *MNRAS*, **488**, 1501
 Grudić M. Y., Guszejnov D., Hopkins P. F., Offner S. S. R., Faucher-Giguère C.-A., 2021, *MNRAS*, **506**, 2199
 1911233, 20009234, 2108318, NSF CAREER grant 1455342, NASA grants 80NSSC18K0562, HST-AR-15800. CAFG was supported by NSF through grants AST-2108230 and AST-2307327; by NASA through grants 21-ATP21-0036 and 23-ATP23-0008; and by STScI through grant JWST-AR-03252.001-A.

Grudić M. Y., Guszejnov D., Offner S. S. R., Rosen A. L., Raju A. N., Faucher-Giguère C.-A., Hopkins P. F., 2022, *MNRAS*, **512**, 216
 Guo M., Stone J. M., Quataert E., Kim C.-G., 2024, arXiv e-prints, p. arXiv:2405.11711
 Guszejnov D., Hopkins P. F., Grudić M. Y., 2018, *MNRAS*, **477**, 5139
 Guszejnov D., Grudić M. Y., Hopkins P. F., Offner S. S. R., Faucher-Giguère C.-A., 2021, *MNRAS*, **502**, 3646
 Guszejnov D., Markey C., Offner S. S. R., Grudić M. Y., Faucher-Giguère C.-A., Rosen A. L., Hopkins P. F., 2022a, *MNRAS*, **515**, 167
 Guszejnov D., Grudić M. Y., Offner S. S. R., Faucher-Giguère C.-A., Hopkins P. F., Rosen A. L., 2022b, *MNRAS*, **515**, 4929
 Guszejnov D., Raju A. N., Offner S. S. R., Grudić M. Y., Faucher-Giguère C.-A., Hopkins P. F., Rosen A. L., 2023, *MNRAS*, **518**, 4693
 Heinemann T., Papaloizou J. C. B., 2009, *MNRAS*, **397**, 52
 Henkel C., Peck A. B., Tarchi A., Nagar N. M., Braatz J. A., Castangia P., Moscadelli L., 2005, *A&A*, **436**, 75
 Hopkins P. F., 2010, arXiv e-prints, arXiv:1009.4702 [astro-ph], Hopkins P. F., 2013, *MNRAS*, **428**, 2840
 Hopkins P. F., 2015, *MNRAS*, **450**, 53
 Hopkins P. F., 2016, *MNRAS*, **462**, 576
 Hopkins P. F., 2017, *MNRAS*, **466**, 3387
 Hopkins P. F., 2023, *MNRAS*, **518**, 5882
 Hopkins P. F., 2024, arXiv e-prints, p. arXiv:2407.00160
 Hopkins P. F., Grudić M. Y., 2019, *MNRAS*, **483**, 4187
 Hopkins P. F., Quataert E., 2010a, *MNRAS*, **405**, L41
 Hopkins P. F., Quataert E., 2010b, *MNRAS*, **407**, 1529
 Hopkins P. F., Quataert E., 2011a, *MNRAS*, **411**, L61
 Hopkins P. F., Quataert E., 2011b, *MNRAS*, **415**, 1027
 Hopkins P. F., Raives M. J., 2016, *MNRAS*, **455**, 51
 Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Martini P., Robertson B., Springel V., 2005a, *ApJ*, **630**, 705
 Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2005b, *ApJ*, **632**, 81
 Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2006a, *ApJS*, **163**, 1
 Hopkins P. F., Hernquist L., Cox T. J., Robertson B., Springel V., 2006b, *ApJS*, **163**, 50
 Hopkins P. F., Hernquist L., Cox T. J., Robertson B., Krause E., 2007, *ApJ*, **669**, 67
 Hopkins P. F., Cox T. J., Kereš D., Hernquist L., 2008, *ApJS*, **175**, 390
 Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, **417**, 950
 Hopkins P. F., Quataert E., Murray N., 2012, *MNRAS*, **421**, 3488
 Hopkins P. F., Torrey P., Faucher-Giguère C.-A., Quataert E., Murray N., 2016, *MNRAS*, **458**, 816
 Hopkins P. F., et al., 2018, *MNRAS*, **480**, 800
 Hopkins P. F., Grudić M. Y., Wetzel A., Kereš D., Faucher-Giguère C.-A., Ma X., Murray N., Butcher N., 2020, *MNRAS*, **491**, 3702
 Hopkins P. F., Squire J., Butsky I. S., 2022a, *MNRAS*, **509**, 3779
 Hopkins P. F., Butsky I. S., Panopoulou G. V., Ji S., Quataert E., Faucher-Giguère C.-A., Kereš D., 2022b, *MNRAS*, **516**, 3470
 Hopkins P. F., Squire J., Butsky I. S., Ji S., 2022c, *MNRAS*, **517**, 5413
 Hopkins P. F., et al., 2023a, *MNRAS*, **519**, 3154
 Hopkins P. F., Butsky I. S., Ji S., Kereš D., 2023b, *MNRAS*, **522**, 2936
 Hopkins P. F., Nadler E. O., Grudić M. Y., Shen X., Sands I., Jiang F., 2023c, *MNRAS*, **525**, 5951
 Hopkins P. F., Grudić M. Y., Kremer K., Offner S. S. R., Guszejnov D., Rosen A. L., 2024a, arXiv e-prints, p. arXiv:2404.08046
 Hopkins P. F., et al., 2024b, *The Open Journal of Astrophysics*, **7**, 18
 Hopkins P. F., et al., 2024c, *The Open Journal of Astrophysics*, **7**, 19
 Hopkins P. F., et al., 2024d, *The Open Journal of Astrophysics*, **7**, 20
 Huang J., Jiang Y.-F., Feng H., Davis S. W., Stone J. M., Middleton M. J., 2023, *ApJ*, **945**, 57
 Jacobs V., Sellwood J. A., 2001, *ApJ*, **555**, L25
 Jafari A., 2019, arXiv e-prints, p. arXiv:1904.09677
 Jiang Y.-F., Dai L., 2024, arXiv e-prints, p. arXiv:2408.16856
 Jiang Y.-F., Cantiello M., Bildsten L., Quataert E., Blaes O., 2015, *ApJ*, **813**, 74
 Jiang Y.-F., Stone J. M., Davis S. W., 2019, *ApJ*, **880**, 67
 Jiménez-Vicente J., Mediavilla E., Kochanek C. S., Muñoz J. A., Motta V., Falco E., Mosquera A. M., 2014, *ApJ*, **783**, 47
 Johansen A., Levin Y., 2008, *A&A*, **490**, 501
 John T. L., 1988, *A&A*, **193**, 189
 Kaaz N., Liska M. T. P., Jacquemin-Ide J., Andelman Z. L., Musoke G., Tchekhovskoy A., Porth O., 2022, arXiv e-prints, p. arXiv:2210.10053

Kaaz N., Liska M., Tchekhovskoy A., Hopkins P. F., Jacquemin-Ide J., 2024, H-AMR FORGE'd in FIRE I: Magnetic state transitions, jet launching and radiative emission in super-Eddington, highly magnetized quasar disks formed from cosmological initial conditions ([arXiv:2410.01877](https://arxiv.org/abs/2410.01877)), <https://arxiv.org/abs/2410.01877>

Kamraj N., et al., 2022, *ApJ*, **927**, 42

Kaspi S., Maoz D., Netzer H., Peterson B. M., Vestergaard M., Jannuzi B. T., 2005, *ApJ*, **629**, 61

King A., 2003, *ApJ*, **596**, L27

Kishimoto M., Antonucci R., Blaes O., Lawrence A., Boisson C., Albrecht M., Leipski C., 2008, *Nature*, **454**, 492

Kormendy J., Bender R., Cornell M. E., 2011, *Nature*, **469**, 374

Krolik J. H., 1999, Active galactic nuclei : from the central black hole to the galactic environment. Princeton, N. J. : Princeton University Press,

Krolik J. H., McKee C. F., Tarter C. B., 1981, *ApJ*, **249**, 422

Kudoh Y., Wada K., Norman C., 2020, *ApJ*, **904**, 9

Kuiper R., Klahr H., Beuther H., Henning T., 2012, *A&A*, **537**, A122

Kurosawa R., Proga D., 2009, *MNRAS*, in press [arXiv:0905.2965],

Laor A., Fiore F., Elvis M., Wilkes B. J., McDowell J. C., 1997, *ApJ*, **477**, 93

Laor A., Barth A. J., Ho L. C., Filippenko A. V., 2006, *ApJ*, **636**, 83

Lenzuni P., Chernoff D. F., Salpeter E. E., 1991, *ApJS*, **76**, 759

Levermore C. D., 1984, *Journal of Quantitative Spectroscopy and Radiative Transfer*, **31**, 149

Lin M.-K., 2015, *MNRAS*, **448**, 3806

Ma X., Hopkins P. F., Wetzel A. R., Kirby E. N., Anglés-Alcázar D., Faucher-Giguère C.-A., Kereš D., Quataert E., 2017, *MNRAS*, **467**, 2430

Magorrian J., et al., 1998, *AJ*, **115**, 2285

Mercedes-Feliz J., et al., 2023, *MNRAS*,

Mihalas D., Mihalas B. W., eds, 1984, Foundations of radiation hydrodynamics. New York, Oxford University Press, 731 p.

Modjaz M., Moran J. M., Kondratko P. T., Greenhill L. J., 2005, *ApJ*, **626**, 104

Murray N., Chiang J., Grossman S. A., Voit G. M., 1995, *ApJ*, **451**, 498

Murray N., Quataert E., Thompson T. A., 2005, *ApJ*, **618**, 569

Narayan R., Yi I., 1995, *ApJ*, **444**, 231

Narayan R., Mahadevan R., Quataert E., 1998, in M. A. Abramowicz, G. Björnsson, & J. E. Pringle ed., Theory of Black Hole Accretion Disks; Cambridge University Press, pp 148–+

Narayan R., Igumenshchev I. V., Abramowicz M. A., 2003, *PASJ*, **55**, L69

Ni Q., et al., 2020, *ApJ*, **889**, L37

Orr M. E., Hayward C. C., Hopkins P. F., 2019, *MNRAS*, **486**, 4724

Ouyed R., Pudritz R. E., 1999, *MNRAS*, **309**, 233

Paczyński B., Wiita P. J., 1980, *A&A*, **88**, 23

Peterson B. M., 1997, An Introduction to Active Galactic Nuclei. Cambridge, New York Cambridge University Press

Peterson B., 2006, The Broad-Line Region in Active Galactic Nuclei. Springer Berlin Heidelberg (Eds. Alloin, Danielle and Johnson, Rachel and Lira, Paulina), Berlin, Heidelberg, pp 77–100, doi:10.1007/3-540-34621-X_3, https://doi.org/10.1007/3-540-34621-X_3

Powell K. G., Roe P. L., Linde T. J., Gombosi T. I., De Zeeuw D. L., 1999, *Journal of Computational Physics*, **154**, 284

Pringle J. E., 1981, *ARA&A*, **19**, 137

Proga D., Stone J. M., Kallman T. R., 2000, *ApJ*, **543**, 686

Ren G., Sun M., Wang J.-X., Cai Z.-Y., 2024, *ApJ*, **967**, 25

Richards G. T., et al., 2006, *ApJS*, **166**, 470

Salpeter E. E., 1964, *ApJ*, **140**, 796

Schmidt M., 1963, *Nature*, **197**, 1040

Semenov D., Henning T., Helling C., Ilgner M., Sedlmayr E., 2003, *A&A*, **410**, 611

Shakura N. I., Sunyaev R. A., 1973, *A&A*, **24**, 337

Shen X., Hopkins P. F., Faucher-Giguère C.-A., Alexander D. M., Richards G. T., Ross N. P., Hickox R. C., 2020, *MNRAS*, **495**, 3252

Shi Y., Kremer K., Grudić M. Y., Gerling-Dunsmore H. J., Hopkins P. F., 2023, *MNRAS*, **518**, 3606

Shi Y., Kremer K., Hopkins P. F., 2024a, *arXiv e-prints*, p. arXiv:2405.12164

Shi Y., Kremer K., Hopkins P. F., 2024b, *ApJ*, **969**, L31

Silk J., Rees M. J., 1998, *A&A*, **331**, L1

Soliman N. H., Hopkins P. F., 2023, *MNRAS*, **525**, 2668

Soltan A., 1982, *MNRAS*, **200**, 115

Spruit H. C., 1996, in Wijers R. A. M. J., Davies M. B., Tout C. A., eds, NATO Advanced Study Institute (ASI) Series C Vol. 477, Evolutionary Processes in Binary Stars. Kluwer academic publishers. Evolutionary processes in binary stars, pp 249–286

Squire J., Quataert E., Hopkins P. F., 2024, *arXiv e-prints*, p. arXiv:2409.05467

Sturm E., et al., 2011, *ApJ*, **733**, L16+

Su K.-Y., Hopkins P. F., Hayward C. C., Faucher-Giguère C.-A., Kereš D., Ma X., Robles V. H., 2017, *MNRAS*, **471**, 144

Su K.-Y., et al., 2021, *MNRAS*, **507**, 175

Torrey P., et al., 2020, *MNRAS*, **497**, 5292

Tortosa A., et al., 2022, *MNRAS*, **509**, 3599

Touma J. R., 2002, *MNRAS*, **333**, 583

Toyouchi D., Inayoshi K., Hosokawa T., Kuiper R., 2021, *ApJ*, **907**, 74

Trakhtenbrot B., et al., 2019, *Nature Astronomy*, **3**, 242

Tremaine S., 2001, *AJ*, **121**, 1776

Utsumi A., Ohsuga K., Takahashi H. R., Asahina Y., 2022, *ApJ*, **935**, 26

Vanden Berk D. E., et al., 2001, *AJ*, **122**, 549

Vlemmings W. H. T., Bignall H. E., Diamond P. J., 2007, *ApJ*, **656**, 198

Walch S., et al., 2015, *MNRAS*, **454**, 238

Weingartner J. C., Draine B. T., 2001, *ApJ*, **548**, 296

Wellons S., et al., 2023, *MNRAS*, **520**, 5394

Williams R. J., Maiolino R., Kronegold Y., Carniani S., Cresci G., Mannucci F., Marconi A., 2017, *MNRAS*, **467**, 3399

Zakamska N. L., et al., 2016, *MNRAS*, **459**, 3144

Zhu S. F., Timlin J. D., Brandt W. N., 2021, *MNRAS*, **505**, 1954

Zier O., Springel V., 2022, *MNRAS*, **517**, 2639

APPENDIX

A. COOLING PHYSICS AND OPACITIES

Our default thermochemistry and cooling physics follows [Grudić et al. \(2021\)](#); [Hopkins et al. \(2023a\)](#) and is described therein, with updates specified here. Briefly, the cooling rates $\Lambda(x, t, T, \rho, Z, x_e, \dots)$ include tabulations and fitting functions for a wide range of processes including detailed pre-tabulated rates for free-free and bound-free transitions (including spontaneous, collisional, Compton, cyclotron/synchrotron, Bremsstrahlung, fine-structure, ionization and recombination) of 11 different separately followed elements (H, He, C, N, O, Ne, Mg, Si, S, Ca, Fe), dust grains (modeling a full size spectrum from nm through $\sim 0.1 \mu\text{m}$ sizes following [Weingartner & Draine 2001](#); [Semenov et al. 2003](#)), free electrons, charged grains, ions (including important negative ions like H^- as well as partially-ionized H, He, C, etc.), atomic and molecular (H_2 , HD, CO, and heavier) species. For all of these the salient abundances of each species are evolved implicitly from the non-equilibrium network and the rates depend on the relative abundances, metallicities, temperatures, etc.

A variety of non-radiative heating/cooling processes are explicitly accounted for as well, including for example cosmic-ray heating (via streaming and hadronic losses and ionization) and ionization; dust collisional and photo-electric heating/cooling; charge exchange reactions; and the “MHD work” from adiabatic gas evolution, shocks, and reconnection (given by the Riemann problem in each cell). The radiative absorption/ionization and cooling mechanisms all couple to the evolved radiation fields but also account for additional radiation sources such as the (appropriately redshift-dependent; [Faucher-Giguère 2020](#)) meta-galactic UV background and CMB (accounting for self-shielding at higher column densities).

The dust-to-metals mass ratio is approximated by $f_{\text{dm}} = 0.5 \mathcal{F}_{\text{sub}}(T_{\text{dust}}/T_{\text{sub}})$ with $T_{\text{sub}} \equiv 1500 \text{ K}$ and $\mathcal{F}_{\text{sub}}(x) \equiv 0.5 (1 + 9(1-x)/\sqrt{1+81(1-x)^2}) \exp[-(x/3)^2]$. We approximate two-temperature (different proton and electron temperatures) plasma effects as detailed in ([Faucher-Giguère & Quataert 2012](#); [Hopkins et al. 2016](#); [Torrey et al. 2020](#)), evolving ion temperatures explicitly limited by Coulomb exchange at sufficiently high temperature. We account for self-absorption of ionizing photons within a cell explicitly (treating each as a slab). For self-absorption of continuum radiative transitions within a cell we implement a simple cor-

rection by multiplying the optically-thin emission $\Lambda_{\text{rad}, i}$ of that cell i by $1/(1 + \tau_{\text{self}, i}^2)$ where $\tau_{\text{self}, i} \equiv \kappa_{\text{rad}, i}(T_{\text{rad, em}} = T_{\text{gas}, i}, \rho_i, T_{\text{dust}, i}, Z_i, \dots) \rho_i \Delta x_i / 2$, i.e. some average continuum self-opacity. But we find that correction is generally small.

Also in our default simulations, we calculate the opacities for each cell and each evolved radiation band based on the combination of the explicitly-integrated chemical network results plus T_{rad} and other evolved radiation properties, including molecular (scaling with CO and H₂ abundances as $\kappa_{\text{mol}} \sim 0.1 (f_{\text{CO}} + 3 \times 10^{-9} f_{\text{H}_2})$, where f_x is a mass fraction and units are all in CGS), dust (scaling with the sublimation-dependent dust-to-metals ratio, dust temperature, T_{rad} , and metallicities according to the detailed tables in Semenov et al. 2003); free electron ($\kappa_e \sim 0.4 x_e (1 + 2.7 \times 10^{11} \rho T_{\text{gas}}^{-2})^{-1} [1 + (T_{\text{rad}}/4.5 \times 10^8)^{0.86}]^{-1}$ where x_e is the free e^- fraction per nucleon, ρ and T_{gas} are the gas density and temperature and T_{rad} the radiation temperature); Kramers continuum ($\kappa_{\text{Kramers}} \sim 4 \times 10^{25} (1 + f_{\text{H}}) (f_Z e^{-1.5 \times 10^5 / T_{\text{rad}}} + 0.001 x_e) \rho T_{\text{rad}}^{-3} T_{\text{gas}}^{-1/2}$ where f_Z is a total metal mass fraction *not* locked into dust); Rayleigh (atomic scattering $\kappa_{\text{Rayleigh}} \sim f_{\text{HI}} \text{MIN}[5 \times 10^{-19} T_{\text{rad}}^4, 0.2 (1 + f_{\text{H}})]$); iron line-blanketing opacities ($\kappa_{\text{Fe}} \sim 1.5 \times 10^{20} f_Z \rho T_{\text{rad}}^{-2} e^{-(8000/T_{\text{rad}})^4 - (T_{\text{rad}}/7 \times 10^5)^2}$, an approximate fit to the tabulated scalings in Jiang et al. 2015 where they are important); and H⁻ ($\kappa_{\text{H}^-} \sim x_{\text{H}^-} (k_{\text{H}^-}^{\text{bf}} + k_{\text{H}^-}^{\text{ff}})$, $k_{\text{H}^-}^{\text{bf}} \sim 4.2 \times 10^7 (8760/T_{\text{rad}})^{3/2} e^{-8760/T_{\text{rad}}}$, $k_{\text{H}^-}^{\text{ff}} \sim 1.9 \times 10^6 (8760/T_{\text{rad}})^2 e^{-8760/T_{\text{rad}}} (0.6 - 2.5 \phi_-^{1/2} + 2.5 \phi_- + 2.7 \phi_-^{3/2})$ with $\phi_- = \text{MIN}(T_{\text{gas}}/5040, 2)$, with the calculation of these and x_{H^-} following John 1988; Lenzuni et al. 1991; Glover & Jappsen 2007), as well as the salient absorption opacities in each narrow band (Hopkins et al. 2023a). These all contribute to absorption opacities except free electron (scattering only) and dust (which contributes to both scattering and absorption per the tables given).

Total energy is manifestly conserved in all radiative absorption/emission/cooling/heating operations. Radiative cooling is added directly into the photon energy budget, and absorption to the gas energy (Grudić et al. 2021). Note that these approaches mean that the gas, dust, radiation, and (implicitly) salient excitation temperatures (determining the relative species/ionization/excitation abundances) are all evolved *separately*, i.e. we do not impose the common simplifying assumption that gas and radiation temperatures must be in equilibrium nor that the system is blackbody-like. Similarly note the appearance of terms like x_e , f_{H} , x_{H^-} , f_{H_2} in the equations above, which are integrated explicitly in our thermochemical network; this means there is no assumption of local thermodynamic equilibrium (LTE), nor of any equilibrium for the slowest-evolving species, in our opacities. Related, our treatment is able to handle both optically thin or highly non-LTE or non-equilibrium limits, and thick or photon-trapped or blackbody limits, without prior assumptions.

However, the above split treatment of opacities and radiative cooling is not always strictly self-consistent, in the sense that the optically-thin radiative cooling emission per particle $n \Lambda_{\text{cool, rad}, i}$ calculated in the cooling operation will not always be exactly equal to $4\pi \langle \kappa_{\text{abs}, i} \rangle \mu_i m_p \sigma_B T_{\text{gas}, i}^4$, the emission rate one would estimate using the band-averaged opacities $\kappa_{\text{abs}, i}$ defined above. This is intentional: since the thermo-chemistry is strictly local for a given incident radiation field we can integrate emission rates over an

arbitrarily fine-grained set of frequencies and species and include non-radiative and non-equilibrium effects per above, but for computational expense the explicitly-evolved radiation fields must be coarse-grained/binned, breaking this exact symmetry. But we can test the approximate consistency of the two with a couple of simple tests. First, for the opacities in our adaptive grey-band (focusing on that as the narrow-band cases are less complex), we could take $\kappa_{\text{abs, eff}, i}(T_{\text{rad}, i}) \rightarrow n \Lambda_{\text{cool, rad}, i}(T'_{\text{gas}}, T_{\text{rad}, i}, \mathbf{x}_i, t_i, n_i, Z_i, \dots) / (4\pi \mu_i m_p \sigma_B T_{\text{gas}}^4)$ where $T'_{\text{gas}} = T_{\text{rad}, i}$, i.e. forcing the mean absorption opacities to match the emission from the thermo-chemical network at the radiation temperature. We stress that this is not the default because it is often less accurate: it ignores the differences between e.g. Planck or Rosseland or flux mean opacities, and can significantly overestimate absorption opacities in certain limits especially relevant in the more diffuse outer disk radii (more ISM-like conditions), for example when the radiation bands are broad-band continuum but the cooling is dominated by recombination transitions or narrow lines, or when the radiation is otherwise highly non-thermal dominated or optically-thin (so the effective radiation temperature of the emission is not, in fact, closely related to the gas kinetic temperature). Second, we can conversely replace our detailed cooling rates with $\Lambda_{\text{cool, rad}, i} \rightarrow 4\pi n_i^{-1} \mu_i m_p \sigma_B T_{\text{gas}}^4 \kappa_{\text{abs, eff}, i}(T'_{\text{rad}}, T_{\text{gas}, i} \mathbf{x}_i, t_i, n_i, Z_i, \dots)$ where $T'_{\text{rad}} = T_{\text{gas}, i}$, i.e. the value given by this our absorption opacity tables at a radiation temperature equal to the gas temperature. Again, this is less accurate for the reasons above. But in the limit where the radiation and gas kinetic temperatures are broadly similar, or the radiation is broad-band continuum-dominated, or the system is sufficiently optically-thick, these should all agree reasonably well (note these experiments still allow for arbitrarily non-LTE behavior, and separately evolve the various abundances and temperatures above). Indeed we find this is the case: in the inner disk of interest here (within a few hundred R_g), these experiments introduce quantitative differences, but do not qualitatively change any of our conclusions or key thermochemical or radiation properties of the system.

B. PHYSICAL AND NUMERICAL VARIATIONS CONSIDERED

We have experimented with a number of alternative physics or numerics experiments and setups, to ensure the behaviors here are robust and to check different potential uncertainties. For brevity, we summarize here the variant runs, restarted at different times from snapshots of our fiducial simulation or parent simulation, and generally run for a much short time (though always still many dynamical times at the inner boundary), to test if they strongly influence the results.

Refinement variations: We considered several different functions to define the “target mass” Δm versus R , including smoother/shallower refinement beginning at larger radii, steeper refinement, or even refining entirely in the initial conditions (akin to simply running a higher-resolution IC statistically sampling the density field as interpolated from the low-resolution snapshot, and hoping there would be no contamination by low-resolution cells migrating inwards). Per § D we also considered pure spatial-resolution refinement criteria (enforcing $\Delta x_i \equiv (m_i/\rho_i)^{1/3} < 0.01 r$). We disabled de-refinement of outflows (any material with $v_r > 1 v_K$ or $0.5 v_K$), or any gas with local thermal-scale lengths c_s/Ω

that would become un-resolved upon de-refinement. We also shifted the “pivot radius” of refinement from $\sim 0.001 - 0.1$ pc. We varied the enforced minimum number of timesteps between allowed refinement or de-refinement operations per cell from $\sim 3 - 3000$ (see Paper I); ran tests in which the temperatures and/or ionization states were arbitrarily reset to $\sim 10^4$ K and/or fully-ionized (sacrificing energy conservation) after a refinement step (to see how quickly the chemistry reverted to standard behavior); forced the code to revert to a first-order reconstruction and solver for any cell which experienced (de)refinement in the last 10 timesteps (see Hopkins 2015); applied an extra 10-iteration Poisson-type divergence cleaning layer following Hopkins (2016) after each refinement; adopted the alternative method from Su et al. (2021) to assign \mathbf{B} after refinement to maintain divergence control; and varied the separation distance and axis of separation of the child cells within each parent cell split (Grudić et al. 2021). So long as the refinement was sufficiently rapid that the target mass Δm could be reached before gas accreted through our boundary, and sufficiently slow so as to not create memory errors or prevent simulation progress (by e.g. creating billions of new cells in a short number of timesteps), we found none of these choices has a systemic effect on our results interior to the high-resolution region.

Resolution variations: Our default refinement pushes to the boundary of what is technically possible, so we cannot simply increase the resolution further. However, we have considered resolution tests in Paper I, and we experiment here with refining further to $\sim 30\times$ improved mass resolution in a single narrow radial annulus. Even that is challenging, as material moves in and out rapidly but making the annulus much larger increases computational expense excessively, so we evolve for just ~ 1 orbital time, but see no qualitative effect on disk structure in this experiment. More easily, we run an experiment where we degrade the target mass resolution everywhere by an order of magnitude. As expected we see less small-scale disk structure and turbulence, but the zeroth-order properties like accretion rates and disk thickness remain similar everywhere (at the very innermost radii, this does make the vertical disk thickness somewhat poorly resolved, however). The refinement variations above also provide resolution tests of different sorts, as several (such as fixed spatial resolution, or disabling de-refinement in winds or cold/short thermal-scale-length gas) provide enhanced resolution in different gas phases (at the cost of more poor resolution in other phases). We describe some additional resolution tests in § D.

Inner boundary variations: We have varied the inner boundary/accretion radius in shorter restarts from $\sim 3 - 300 R_g$. The larger radii approach our original simulation accretion radius, so are consistent with those simulations as expected. But the results *at a given radius* do not appear strongly sensitive to this choice. The main limitation of going to smaller radii is the much smaller dynamical and light-crossing times (and therefore numerical timesteps) involved, limiting the duration of the experiments. Below we show the structure that appears in our experiments with smaller inner radii.

Initial conditions: We restarted our simulation beginning at three different times (snapshots) randomly sampled from the parent simulation after it had reached quasi-steady state (see Paper II). We also consider snapshots from both the “fast” and “slow” protostellar disk accretion sub-grid model variant simulations in Hopkins et al. (2024a) – this choice only changes the

sink-particle description of accretion for stars forming in the accretion disk, so unsurprisingly has no effects here. We also experimented with altering our outer “frozen” radius, making it larger or smaller by a factor of ~ 3 . We ran tests in which we reset all initial temperature and/or ionization states of the snapshot before restart (making the ICs uniformly $T = 10^4$ K and/or fully ionized). We also considered a run in which the inner boundary gradually shrunk (decreasing by a factor of ~ 2 every two orbital times at each radius) as opposed to simply setting it to $\sim 30 R_g$ initially. Again none of these had a significant impact on our results. As emphasized in the text, because our simulations are trying to predict uncertain quantities like the inner disk accretion rates, luminosities, and outflow properties, our inner boundary is a simple sink, but we have also considered simulations where the inner boundary features an outgoing steady radiation flux with a luminosity $= 3 \times 10^{45}$ erg s $^{-1}$ directed with a $|\cos \theta|$ intensity pattern outwards at the inner boundary; this can have some significant effects on the inner boundaries in simulations where the inner boundary is large ($\sim 300 R_g$ or larger), but once we refine to small radii, we find that starting from this modified IC makes no difference to the qualitative disk structure that forms (it is similar to restarting from a different snapshot in our default setup).

Star formation and stellar physics: As noted below, we experimented with removing any star particles, turning off new star formation and stellar physics entirely, or considering the different fast/slow sub-grid circumstellar accretion models from Paper III. None of these has any effect at radii $\sim 3 - 300 R_g$ where we focus here, as there are basically no stars and no star formation in any of the cases considered.

Gravitational potentials: We experimented with switching from a Paczyński-Wiita (PW) potential for the BH back to Newtonian, replacing the gas cells potential with PW (though not self-consistent, this is irrelevant because their individual cell R_g is vastly smaller than our resolution), or disabling entirely the external potential from matter exterior to the “frozen” larger radii \gtrsim pc. None of these had any major effect at $\gtrsim 30 R_g$, and even removing the PW potential at $\sim 3 - 30 R_g$ only produces modest quantitative (not qualitative) effects on the disk properties.

Radiation-hydrodynamics: We have considered removing all radiation transport, either assuming radiation is locked to gas (so contributes to the pressure but always has zero cell-to-cell flux, akin to an adiabatic-like $\gamma = 4/3$ equation-of-state assumption) or simply deleting all radiation immediately (assuming optically-thin escape of all cooling radiation), which both have significant effects as described in the main text. Less radically, we can also turn on and off different radiation bands of our multi-group solver in a modular fashion. In the inner disk, we can remove the EUV, NUV, and OIR bands and reduce to a single photo-ionizing band and see very little effect, since at these wavelengths narrow-band effects like Lyman-Werner radiative dissociation or photo-electric heating are not the dominant sources of radiation or opacity. Removing any treatment of ionizing photons (e.g. removing all ionizing bands and not calculating ionization from the adaptive-wavelength band) produces an artificial result that the disk remains atomic at all radii, which in turn means the opacities are orders-of-magnitude lower than they should be, and so the effective optical depth $\tau^* \sim \sqrt{\tau_a(\tau_a + \tau_s)}$ remains below unity and the disk never thermalizes or effectively traps radiation as it should. But so long as this is accounted for, we can capture the

most important physics with our single adaptive band as this still allows for a flexible scaling of T_{rad} which is very different across the disk. We have modified various numerical slope and flux-limiters used to stabilize the operator-splitting of the thermochemistry and radiation transport (see Hopkins 2017), and experimented with the alternative operator-splitting approach used in GIZMO radiation-hydrodynamic TDE simulations in Bonnerot et al. (2021), and switched from our default comoving formulation of the RHD equations (since the code is Lagrangian) to a boosted-Eulerian mixed-frame RHD formulation per Mihalas & Mihalas (1984). None of these technical numerical choices produces significant qualitative effects.

Cooling physics & Opacities: More details of the cooling physics and opacities included are given in § A. We have experimented with a number of the salient choices there. For example, turning off the self-absorption correction for continuum radiation, relativistic Compton cooling (just using non-relativistic expressions) and Comptonization, including or excluding Klein-Nishina corrections to free electron opacity, excluding our approximate treatment of two-temperature plasmas (relevant for extremely hot gas), removing the iron line-blanketing opacities, changing the dust sublimation temperature (and rapidity of sublimation above T_{sub}) between $\sim 1000\text{--}2000$ K, or adding all cooling radiation to our adaptive continuum band versus spreading it between bins according to the estimated emission wavelengths. Most of these have at most modest (order-unity) quantitative effects on gas temperatures in some specific regions of parameter space, but do not significantly alter the mean (mass-weighted) gas temperatures or ionization structure or radiation temperatures, let alone the global disk dynamics. The same is true if we replace our default cooling network with just our grey-band opacities or conversely calculate the effective opacities not directly but simply from the cooling subroutine emissivities, as described in § A – these can produce significant inaccuracy in the outer, optically-thin atomic disk (more ISM-like conditions) but for the innermost radiation-pressure dominated disk more simple approximations appear to behave reasonably.

C. ON THE “MINIMUM RESOLVABLE” SCALE HEIGHTS/LENGTHS

In Eulerian (e.g. fixed-grid) MHD methods, the minimum spatial scale $\Delta x_{\text{min}} = \Delta x$ which can be represented is trivially defined by the cell size, but there is no unique definition of the minimum mass scale $\Delta m_{\text{min}} = \rho_{\text{min}} \Delta x^3$ (in terms of some “minimum possible density” ρ_{min}). Conversely, with a Lagrangian methods (like ours), the minimum mass scale $\Delta m_{\text{min}} = \Delta m$ is fixed, but there is no unique definition of the minimum resolvable spatial scale $\Delta x_{\text{min}} = (\Delta m / \rho_{\text{max}})^{1/3}$ (noting the 1D cubic-cell equivalent cell size is $(m_i / \rho_i)^{1/3}$, and now in terms of a “maximum possible density” ρ_{max}). Some examples of the latter salient for different physical applications for our GIZMO code (e.g. Jeans fragmentation) are discussed in Hopkins (2013, 2015); Hopkins & Raives (2016); Hopkins et al. (2018).

A well-posed and salient question here is what the minimum resolvable disk thickness H_{min} would be. If we started, for example, from an arbitrarily tall gas cylinder with $H \rightarrow \infty$, in a stratified potential, with only thermal pressure support, and allowed it to cool to zero temperature, it would collapse vertically, with the instantaneous minimum cell size $\text{MIN}(m_i / \rho_i[t])^{1/3}$ shrinking as the disk collapses (so this instantaneous value is clearly not the minimum resolvable).

However, at some point, if the cell centers-of-mass (mesh-generating points, moving with the fluid) collapse to a razor-thin configuration, then the separation in the horizontal (xy or $R\phi$) direction Δx between cells will become much larger than their vertical separation (Δz). At this point, quantities like vertical gradients can no longer be defined (or more formally in GIZMO, the condition matrix for the face and gradient reconstruction becomes ill-conditioned; see Hopkins 2015 Appendix C). This is what it means for a structure to be “one cell thick” in such a method. The vertical thickness at which this occurs is easily calculated as $\Delta z_{\text{min}} \approx (dN/dA)^{-1/2} = (\Delta m / \Sigma_{\text{gas}})^{1/2}$, in terms of the number of cells per unit area (vertically integrating through the disk scale height) dN/dA . Direct numerical tests also validate this scaling for practical questions (Hopkins et al. 2011, 2012; Hopkins 2015; Few et al. 2016; Deng et al. 2017, 2021). This defines the resolution limit for H/R plotted in Fig. 5. Note that this is, by approximate design of our refinement criterion, order-of-magnitude similar to the thermal scale length c_s/v_K of the midplane gas at most radii.

However, this does not mean that local structures with even smaller sizes ($\Delta x \ll \Delta z_{\text{min}}$) cannot be represented. Indeed we see some dense structures in Paper I and here with sizes below this limit. The key is that these are small-scale structures ($\Delta x \ll H$) – a “clump” of gas cells could, for example, collapse under self-gravity in principle in this numerical method to arbitrarily small scales (i.e. form a star). In practice this occurs at large radii in the disk and is truncated only by the sink particle prescription for star formation (see detailed discussion in Hopkins et al. 2024a, but in brief we see at least some structures with sizes ~ 10 au at $R \gtrsim 5$ pc, or $\Delta x/R \lesssim 10^{-5}$). More relevant in the inner disk, strong radiative shocks can generate denser structures, for which Δx can drop to smaller values.

D. TESTS WITH ENFORCED-RESOLVED THERMAL SCALE-LENGTHS

In our simulations, spatial scales like the turbulent energy range ($\sim |\delta v|/\Omega \sim H$), density scale-height (H), Alfvén scale ($\sim v_A/\Omega \sim H$), mean-field and density gradient scale-lengths, as well as dynamical and light-crossing times and mass scales of the disk are well-resolved (Figs. 3, 5). Indeed our typical $\Delta x/H \sim 0.01$ is comparable to the highest resolution often achieved even in historical shearing-box type simulations (see discussion in Zier & Springel 2022; Dhang et al. 2024). But it is more challenging to even uniquely define (let alone resolve) the thermal scale-length $H_{\text{therm}} \sim c_s/\Omega$. In Fig. 5 we see our resolution is comparable to this ($\Delta x \sim H_{\text{therm}}$) defined for a weighted mean c_s in the midplane. But recall the disks are multi-phase with a wide range of T_{gas} and c_s (Fig. 5), so H_{therm} of the hot gas is quite well-resolved, while (especially in the outer disk) some gas cools to molecular temperatures $\lesssim 100$ K giving local H_{therm}/R as small as $\sim 10^{-5}$. Resolving this is not possible even in shearing-box simulations. But this gas is precisely that in which thermal pressure is least significant (they have magnetic, turbulent, and radiation pressures orders-of-magnitude larger than thermal), so we expect finite resolution not to matter for global disk properties (though it could be important for microphysics like the sizes of the smallest “cloudlets” in the BLR), as commonly seen in other simulations of multi-phase disks in galaxy and star formation (Hopkins et al. 2011, 2012, 2018; Walch et al. 2015; Benincasa et al. 2016, 2020; Guszejnov et al. 2018; Grudić et al. 2019,

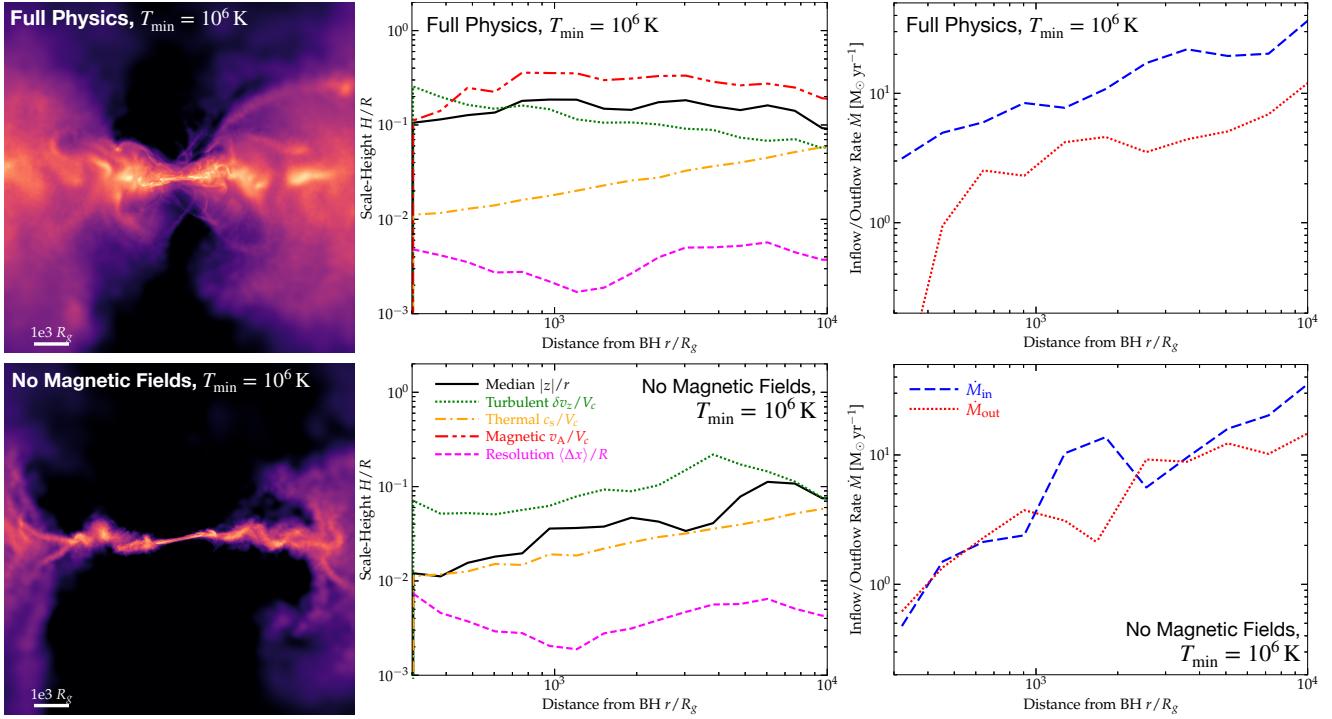


FIG. 31.— Tests from § D where we enforce a well-resolved minimum thermal scale-length in the disk. We re-run a portion of the disk from our simulation snapshots (with inner boundary at $300 R_g$), run for time $\sim 300\Omega^{-1}$ ($\sim 1\Omega^{-1}$) at radius ($10^4 R_g$), adding an arbitrary temperature floor $T > T_{\min} \equiv 10^6$ K in the thermochemistry to ensure the thermal scale-length $H_{\text{therm}} \sim c_s/\Omega$ is well-resolved at all radii shown. We compare full-physics otherwise (top) and a restart where we remove magnetic fields (radiation-hydrodynamics only, bottom). *Left:* Edge-on density projection in a slice through the angular momentum axis as Fig. 2. *Middle:* Density scale height ($|z|/r$) and predicted thermal, turbulent, magnetic scale-heights and resolution in annuli, versus radius r . *Right:* Mass inflow \dot{M}_{in} and outflow \dot{M}_{out} rates through each annulus. With full physics, the un-ambiguously resolved $H_{\text{therm}} \gg \Delta x$ from this test changes some of the small-scale gas clumping, but none of the global properties of the disk: notably the Alfvén, turbulent, and density scale-heights, and accretion rates, remain similar to our default case. Removing magnetic fields however leads to immediate collapse of the disk to the thermal scale-height, with order-of-magnitude lower inflow rates (and even those are mostly radial “sloshing”/spiral/eccentric motions, hence $\dot{M}_{\text{in}} \approx \dot{M}_{\text{out}}$, with the vertical motions δv_z dominated by the prominent warps/breaks seen in the disk).

2022; Orr et al. 2019).

That said, in Squire et al. (2024) and Guo et al., in prep., we have studied the highly-idealized problems in magnetically-dominated disks, specifically shearing boxes and homogeneous collapse, respectively, with greatly simplified physics (isothermal, ideal MHD, in analytic Keplerian potentials, idealized geometries, and no other physics) and different code (ATHENA). While some of the conclusions for $\beta_{\text{therm}} \ll 1$ (though still $\beta_{\text{therm}} \sim 0.01$, much higher than the outer disks here) are similar to our “full-physics” simulations here, in these idealized experiments there is a hint of some non-linear resolution dependence where magnetic flux-loss becomes more efficient when Δx becomes smaller than H_{therm} (or sufficiently small relative to v_A/Ω), before becoming less efficient again at smaller Δx , which can increase β_{therm} to $\sim 0.1 - 1$ in the midplane. Repeating some of these experiments in GIZMO (Tomar et al., in prep.) shows qualitatively similar results, despite the very different method and meaning of resolution (§ C).

In our “full physics” simulations, however, we do not see any obvious resolution dependence of the mean field or disk height H/R (§ B), varying Δx by order-of-magnitude. But again, it is impossible to reach $\Delta x \ll H_{\text{therm}}$ for all cold gas. Therefore here we consider a simple experiment, restarting our fiducial simulation enforcing a temperature floor in the cooling function. We consider both $T_{\min} = 10^6$ K (so $H_{\text{therm}}/R \propto R^{1/2}$) and $T_{\min} \propto R$ (chosen so $H_{\text{therm}}/R = 0.1$), evolving the outer disk (inner boundary moved out to $\sim 300 R_g$) for ~ 20 orbital times

($> 100\Omega^{-1}$) at these radii. These values of T_{\min} are chosen so that the disk retains $\beta_{\text{therm}} \ll 1$, but with a well-resolved hierarchy $v_A/\Omega \gg c_s/\Omega \gg \Delta x$. We also add the refinement criterion that no cell can have individual instantaneous size $\Delta x_i \equiv m_i/\rho_i$ larger than $0.01 r$.⁹ The resulting H/R and vertical profiles are shown in Fig. 31. We see that there is no major loss of toroidal magnetic flux, or change in either the magnetic (v_A/Ω) or turbulent ($\delta v_z/\Omega$) or density (H) scale-heights of the disk. Indeed, considering essentially all the plots from the main text, the only noticeable change is some suppression of small scale structure with sizes $\ll H$ in the disk (e.g. some of the narrow shocks and dense gas clumps), which is completely expected given the much larger thermal scale length and lower (by construction) sonic Mach number.

Thus there is no major qualitative change in global properties of these disks owing to explicitly resolving the thermal scale-length (consistent with our initial speculation), nor is there a change with resolution in the range $\Delta x \sim 0.01 - 1 H$. This suggests that there is some important difference in our global, galaxy-scale, full-physics simulations compared to the idealized tests in Squire et al. (2024) and Guo et al., in prep. This may owe to the effects of finite radiation pressure/live radiation transport, more realistic thermochemistry and multi-phase structure, boundary/initial conditions, pre-existing tur-

⁹ Note, in our tests in § D, the thermal scale-height is well-resolved with $c_s/\Omega \gg \Delta x$ whether we use the definition of the minimum-resolvable Δx from § C as appropriate for a Lagrangian code, or simply take the average $\langle \Delta x_i \rangle \equiv \langle m_i/\rho_i \rangle$ of cells in the disk.

bulence and strong fields, or other effects not captured in those test problems.

If, however, we remove magnetic fields from this experiment in a single timestep (keeping the radiation and turbulent and thermal support fixed), shown in Fig. 31, we can confirm the expected behavior, that the disks collapse rapidly to the input thermal scale-height, without any of its thick vertical structure. But we also see the behaviors noted in Paper II at larger radii if we begin our refinement from a galaxy-scale setup without magnetic fields: the accretion disk actually becomes a *decretion* disk on average, with a widening central hole fea-

turing only occasional very small accretion of gas clumps so that the time-averaged gas accretion rate supported drops by an order of magnitude in a short time (and is continuing to drop at an increasing rate at the final times to which we run the test). In these simulations, magnetic fields play a critical role regardless of how well one resolves the thermal scale-length.

This paper was built using the Open Journal of Astrophysics L^AT_EX template. The OJA is a journal which provides fast and easy peer review for new papers in the astro-ph section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.